

AD-A085 394

RAND CORP SANTA MONICA CA

F/6 5/1

RULE-BASED MODELING AS AN ANALYSIS TOOL: IMPLICATIONS FOR RESOU--ETC(U)

APR 80 R FALLON


F49620-77-C-0023

UNCLASSIFIED

RAND/N-1489-AF

NL

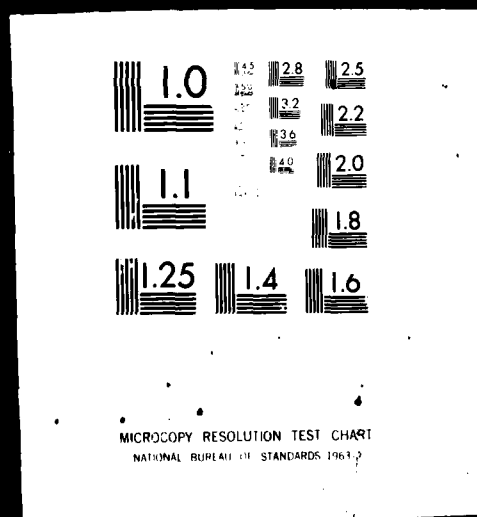
1 OF 2
AD
AD-A085 394



1 OF 2

AD

A085394



ADA 085394

A RAND NOTE

LEVEL II

12
B.S.

RULE-BASED MODELING AS AN ANALYSIS TOOL:
IMPLICATIONS FOR RESOURCE ALLOCATION WITHIN
THE STRATEGIC AIR COMMAND

Richard Fallon

April 1980

N-1489-AF

DTIC
COLLECTED
JUN 12 1980
D
C

Prepared For

The United States Air Force

This document has been approved
for public release and sale; its
distribution is unlimited.

Rand
SANTA MONICA, CA 90406

80 6 10 014

DDC FILE COPY

The research reported here was sponsored by the Directorate of Operational Requirements, Deputy Chief of Staff/Research, Development, and Acquisition, Hq USAF, under Contract F49620-77-C-0023. The United States Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright notation hereon.

The Rand Publications Series: The Report is the principal publication documenting and transmitting Rand's major research findings and final research results. The Rand Note reports other outputs of sponsored research for general distribution. Publications of The Rand Corporation do not necessarily reflect the opinions or policies of the sponsors of Rand research.

Published by The Rand Corporation

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER N-1489-AF	2. GOVT ACCESSION NO. AD-A085394	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Rule-Based Modeling as an Analysis Tool: Implications for Resource Allocation Within the Strategic Air Command.		5. TYPE OF REPORT & PERIOD COVERED Interim rept.
7. AUTHOR(s) R. Fallon Richard Fallon		6. PERFORMING ORG. REPORT NUMBER
8. CONTRACT OR GRANT NUMBER(s) F49620-77-C-0023		9. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
10. CONTROLLING OFFICE NAME AND ADDRESS The Rand Corporation 1700 Main Street Santa Monica, Ca. 90401		11. REPORT DATE Apr 1980
11. CONTROLLING OFFICE NAME AND ADDRESS Requirements, Programs, & Studies Group (AF/RDM) Ofc, DCS/R&D & Acquisition Hq USAF, Washington, D. C. 20330		12. NUMBER OF PAGES 171
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) 12/19/80		13. SECURITY CLASS. (of this report) UNCLASSIFIED
15. DECLASSIFICATION/DOWNGRADING SCHEDULE		
16. DISTRIBUTION STATEMENT (of this Report) Approved for Public Release; Distribution Unlimited		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report) No Restrictions 14/RAND/N-1489-AF		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Decision Making Computer Systems Programs Allocation Models DOSS Scheduling Bomber Aircraft Military Training Flight Crews		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) See Reverse Side 296600 Lm		

UNCLASSIFIED

CLASSIFICATION OF THIS PAGE(When Data Entered)

Examines the potential of rule-based modeling as an analysis tool for investigating resource allocation policy issues. Focus is on resource allocation within B-52 flying organizations of the Strategic Air Command. A rule-based computer system, DOSS (Decision Oriented Scheduling System), is demonstrated to provide a valid model of many variables that affect resource allocation of aircrews and aircraft. DOSS is then used to analyze effects on wing performance of several alternative decision rules and policies. Analysis focuses on the capability of SAC bomb wings, given current resources, to increase the SAC alert force and to fly more training sorties. The analysis has implications for the particular policy issues examined, but its more general aim is to illustrate the potential of DOSS for examining a broad range of resource allocation policy issues within SAC. A number of implications are discussed regarding the potential of rule-based modeling for improving organizational decisionmaking in general. 171 pp. Bibliog. (Author)

SECURITY CLASSIFICATION OF THIS PAGE(When Data Entered)

A RAND NOTE

RULE-BASED MODELING AS AN ANALYSIS TOOL:
IMPLICATIONS FOR RESOURCE ALLOCATION WITHIN
THE STRATEGIC AIR COMMAND

Richard Fallon

April 1980

N-1489-AF

Prepared For

The United States Air Force

DTIC
ELECTE
JUN 12 1980
C



Accession For	
NTIS GRA&I	<input checked="checked" type="checkbox"/>
DDC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By _____	
Distribution/	
Availability Codes	
Dist.	Avail and/or special
A	

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

PREFACE

This study examines the potential of a rule-based computer system called DOSS (Decision Oriented Scheduling System), developed at The Rand Corporation, as an analysis tool for investigating resource allocation policy issues within the Strategic Air Command (SAC). This research was performed under the Project AIR FORCE project "Scheduling and Resource Allocation" for the Directorate of Operations and Readiness, Hq USAF (AF/XOO).

DOSS was developed by Rand to model the scheduling of aircrew and aircraft resources within SAC. In September 1978, after discussions with members of the Hq SAC senior staff, including the Chief of Staff, DCS/Operations, and DCS/Logistics, Rand analysts were asked to examine the potential of DOSS as an analysis tool in the context of the B-52 Aircrew Continuation Training Test Program. DOSS was used first to model SAC wing scheduling (both operations and maintenance) under the new training program, then to examine particular resource allocation issues within SAC. The analysis focused on determining the capability within SAC bomber wings, given current resources, to (a) fly more training sorties and (b) increase the SAC alert force.

Interim results of this research were briefed to Major General Jack L. Watkins, Deputy Chief of Staff for Operations, at Hq SAC in March 1979. For the contents of this briefing, and an overview of the research presented in this report, see Rand Note N-1161-1-AF, DOSS as an Analysis Tool: Applications to the B-52 Aircrew Continuation Training Test Program--A Progress Report, by R. Fallon, July 1979.

This report is being submitted simultaneously to The Rand Graduate Institute as a dissertation in partial fulfillment of the requirements for a doctorate in policy analysis.

SUMMARY

This study examines the potential of rule-based modeling as an analysis tool for investigating resource allocation policy issues within a large public-sector organization. The focus is on aircrew and aircraft scheduling within B-52 flying organizations of the Strategic Air Command (SAC).

SAC's basic goal for manned aircraft is to maintain mission-ready aircrews and aircraft as a credible nuclear deterrent. SAC operates several decentralized bombardment wings, each of which conducts a flying program to maintain proficiency of aircrews and to exercise and maintain reliable ready aircraft. Additionally, some of the aircrew and aircraft resources are continuously assigned to ground alert to serve as a quick-reaction strategic force.

Rule-based modeling of organizational decisionmaking is based on the assumption that decisionmaking within organizations can be characterized by an extensive set of interrelated heuristics or decision rules, which if programmed into a computer can provide valid models of decisionmaking behavior. In this study, a rule-based computer system called DOSS (Decision Oriented Scheduling System) was demonstrated to provide a valid model of the many variables (decision rules, higher-headquarters policies, and resource constraints) that affect resource allocation of aircrews and aircraft at a typical SAC bomb wing.

DOSS was used to analyze the effects on wing performance of several alternative decision rules and policies for allocating resources. The analysis focused on the capability of SAC bomb wings,

given current resources, to (a) increase the SAC alert force, and to (b) fly more training sorties. The analysis has implications for the particular policy issues that were examined, but its more general purpose was to illustrate the potential of DOSS for examining a broad range of resource allocation policy issues within SAC, and possibly within other Air Force Commands as well.

DESIGN OF THE ANALYSIS

The analysis that was undertaken with DOSS focused on rule and policy changes to aircrew and aircraft scheduling, which were designed primarily to increase aircrew and aircraft availability to fly. With regard to aircrew scheduling, an additional policy change was examined that allowed differential training of aircrews. Differential training refers to distributing training sorties to aircrews based on a measure of the relative proficiency of aircrews. This is in contrast to the current policy of giving every crew the same number of sorties.

The effects of these changes, on several measures of wing performance, were analyzed at each of three different sortie and alert levels: (1) the current level of sorties and alert, (2) the current level of sorties but with increased alert, and (3) increased sortie and alert levels. This experimental approach was chosen in order to estimate the feasibility and the effects of achieving (a) increased sortie levels, (b) increased alert levels, and (c) significant differential training, both with and without rule and policy changes.

MAJOR FINDINGS OF THE ANALYSIS

A major finding of the analysis was that increased levels of sorties and alert are possible given current aircrew resources, provided that certain rule and policy changes to increase aircrew availability are adopted (for example, allowing the use of shortened mission planning prior to all sorties). In one case, DOSS scheduled a 12.5 percent increase in total sorties for a three-month period, along with increasing the alert force from four to five aircrews on alert.

Furthermore, the analysis demonstrated that a significant potential for differential training exists at all the sortie and alert levels that were examined. The major implication of differential training is that greater overall training effectiveness can be achieved given any fixed number of sorties.

Besides the positive implications for training and alert that result from (a) more sorties, (b) increased alert, and (c) significant differential training, the analysis also pointed out some negative aspects of increasing the level of sorties and alert. The most serious of these aspects are (a) a possible reduction in the diversity of training sorties, and (b) an increase in the average workweek for aircrews.

Even though the analysis of rule and policy changes to aircrew scheduling demonstrated that additional sorties and alert are possible given current aircrew resources, no analysis of overall wing capability would be complete without considering the effects on maintenance of increasing sortie and alert outputs.

The major finding of the analysis pertaining to aircraft scheduling and maintenance was that aircraft availability is not a

limiting factor in scheduling more sorties and alert. DOSS scheduled both increased sorties as well as increased aircraft on alert, without the need of any policy changes to increase aircraft availability. These increases in output, however, necessitate scheduling more turnaround sorties, which allow a fixed number of aircraft to fly more sorties. Furthermore, increased numbers of turnaround sorties result in less aircraft hours spent in pre-flight and post-flight inspections, which consequently reduces maintenance costs in terms of resources expended in these activities.

It should be stressed, however, that more turnarounds and less pre-flight and post-flight inspections may have adverse effects on aircraft reliability over time. Other means of analysis besides DOSS would have to be employed to estimate the effects on aircraft reliability that may result from changes in aircraft scheduling.

GENERAL IMPLICATIONS

The last section of this report is devoted to a discussion of a broad range of other resource allocation policy issues facing SAC now, or possibly in the future, where DOSS may be useful as an analysis tool. Furthermore, some implications not yet fully explored are discussed for improving resource allocation within SAC by using DOSS in other ways, such as: (a) a training device for wing schedulers, and (b) as an aid in wing-level scheduling.

Finally, a number of suggestions stemming from this research are made, regarding implications of rule-based modeling for improving organizational decisionmaking in general. Basically, rule-based models have potential implications for improving organizational

decisionmaking through their use as (1) descriptive models that simulate decisionmaking within organizations, and (2) as analytic tools for examining alternatives. As analytic tools, rule-based models could be used to either (a) analyze decisionmaking within organizations in order to seek improvements, or (b) analyze decisionmaking of other organizations for the purposes of prediction or control.

The use of rule-based models to analyze decisionmaking within organizations could lead to improved decisionmaking at different levels within an organization. Some suggested implications are that rule-based models could lead to (1) increased search of alternatives, (2) more rapid recognition of potential problems with current decision rules, (3) possible reductions in excess resources paid to avoid uncertain outcomes, (4) increased coordination of decisionmaking across subunits of the organization, and (5) improved analysis of corporate- or headquarters-level policy alternatives. The development and use of DOSS in this study provides some evidence in support of these suggested implications. However, more research is needed to validate some of the generalizations that are made.

The use of rule-based models of other organizations for the purposes of prediction or control is particularly relevant to most governmental regulatory agencies, who seek to influence organizational decisionmaking for some societal purpose. Whether public policies, such as economic or environmental regulations, will lead to their desired ends, depends largely upon how they affect the internal decision processes of the organizations they are directed towards. Rule-based models, to the extent that they can help clarify how

organizations deal internally with policy variables being modified externally, have potential to help depict the likely consequences of alternative public policies. The use of rule-based models as analytic tools in this regard seems worthy of a great deal of further research in a variety of organizational contexts.

ACKNOWLEDGMENTS

This work might never have been undertaken had it not been for the insight of Major General Jack L. Watkins, then Deputy Chief of Staff for Operations, Strategic Air Command, that the B-52 Aircrew Continuation Training Test Program provided a good context within which to investigate the utility of the Decision Oriented Scheduling System (DOSS) as an analysis tool. I am also grateful to a number of staff officers at Hq SAC for their support throughout the research effort. I particularly wish to express my appreciation to the many SAC officers and NCOs who spent many hours explaining the complexities of scheduling flying and maintenance activities.

I am grateful to numerous Rand colleagues, but I am particularly indebted to Morton B. Berman for his patient guidance and constructive criticisms of my work. I am also indebted to Albert P. Williams and David S. C. Chu for their insightful comments and thoughtful suggestions regarding this work. I am also grateful to I. K. Cohen and Steven M. Drezner for their guidance and support of this research. I would also like to acknowledge Richard J. Hillestad whose earlier developmental work on DOSS made this work possible. I gratefully acknowledge the fast, accurate, and indefatigable word processing support of Susan D. Jackson in preparing the final manuscript, and the secretarial support of Donna Saenz throughout this research effort.

CONTENTS

PREFACE.....	iii
SUMMARY.....	v
ACKNOWLEDGMENTS.....	xi
Section	
I. INTRODUCTION.....	1
II. THEORETICAL FRAMEWORK.....	6
Decision Theory and Rule-Based Modeling.....	6
Resource Allocation Within SAC Bomb Wings.....	14
Design Concept of DOSS.....	20
Notes--Section II.....	23
III. MODEL DESCRIPTION AND VALIDATION.....	25
Model Description.....	25
Model Validation.....	35
Notes--Section III.....	39
IV. DESCRIPTION OF THE ANALYSIS.....	42
Scope of the Analysis.....	42
Changes to the DOSS Baseline Model.....	49
Notes--Section IV.....	61
V. ANALYSIS AT THE BASELINE SORTIE AND ALERT LEVEL.....	63
Effects of Changes to Aircrew Scheduling.....	63
Effects of Changes to Aircraft Scheduling.....	79
Summary.....	88
Notes--Section V.....	89
VI. ANALYSIS AT INCREASED ALERT.....	93
Effects of Changes to Aircrew Scheduling.....	94
Effects of Changes to Aircraft Scheduling.....	99
Summary.....	104
Notes--Section VI.....	105
VII. ANALYSIS AT INCREASED SORTIE AND ALERT LEVELS.....	106
Effects of Changes to Aircrew Scheduling.....	107
Effects of Changes to Aircraft Scheduling.....	117
Summary.....	122
Notes--Section VII.....	123
VIII. IMPLEMENTATION ISSUES FOR SCHEDULING.....	125

IX. CONCLUSIONS AND IMPLICATIONS.....	134
Major Findings and Policy Implications of the Analysis..	135
Implications of DOSS for Improving Resource	
Allocation within SAC.....	137
Implications of Rule-Based Modeling for Improving	
Organizational Decisionmaking.....	146
Appendix	
A. TRAINING ITEMS.....	159
B. DISTRIBUTION OF SORTIES.....	163
GLOSSARY.....	168
BIBLIOGRAPHY.....	171

FIGURES

1. MR crew scheduled sortie distributions..... 66
2. Potential for differential training at the sortie level... 71

I. INTRODUCTION

The efficient allocation of resources within large complex organizations is a difficult problem involving decisionmaking at all levels. Public-sector organizations, for the most part, face even more difficult problems in the quest for efficiency than do private-sector organizations. This is due in large part to two related matters: (1) the fact that these institutions are not subject to market forces as are competitive profit-making organizations; and (2) the societal goals which large governmental organizations are asked to serve are often vague and hard to operationalize. Consider, for instance, the goal of our armed forces to maintain a state of "readiness" in order to achieve a "credible" deterrent to war. The existence of such vague and hard-to-operationalize goals in the public sector complicates the management of scarce resources to efficiently achieve these goals.

Given the existence of public-sector organizations, the policy problem facing analysts is one of finding ways to motivate these institutions to efficiently allocate the scarce resources entrusted to them for the purpose of fulfilling societal goals. Since the allocation of scarce resources within organizations is essentially a decisionmaking process, it is thereby essential for the policy analyst to understand and take into account how decisions are made within organizations he is trying to affect.

Over the years there has emerged a body of literature in the field of organizational behavior that deals primarily with decision processes within organizations and their effects on organizational

performance. The basis of this school of thought is the assumption of intendedly rational decisionmaking, but modified to take into account the limits of human cognition which prohibit solving a complex problem in a comprehensively rational way. Because of cognitive limitations, decisionmakers facing complex problems rely, among other things, on a myriad of simplifying decision rules in order to seek solutions. Within an organization such decision rules may take the form of stated policy, but many are unstated rules of thumb or standard operating procedures which are created and changed in an evolutionary process.

Related work in the field of artificial intelligence attempts to simulate, with the aid of the computer, human problem-solving behavior in a variety of contexts. Rule-based modeling (or heuristic modeling, as it sometimes is referred to) is a methodology based on the assumption that many types of complex human problem-solving can be broken down into an extensive set of interrelated heuristics or rules of thumb, which if programmed into a computer can provide valid models of that behavior. Rule-based modeling, when applied to decisionmaking within organizations, is therefore based on assumptions highly consistent with the theoretical view of organizational decisionmaking discussed above. Furthermore, rule-based models, designed to incorporate organizational decision rules, hold promise as analytic tools for examining the effects on organizational performance of alternative sets of rules.

The primary aim of this study is to examine the potential of rule-based modeling as an analysis tool for investigating resource allocation policy issues within a large public-sector organization. More specifically, this investigation focuses upon resource allocation within B-52 flying organizations of the Strategic Air Command (SAC).

SAC's basic goal for manned aircraft is to maintain mission-ready aircrews and aircraft as a credible nuclear deterrent. SAC operates 19 decentralized bombardment wings within the United States, each of which conducts a flying program to maintain proficiency of aircrews and to exercise and maintain reliable ready aircraft. Additionally, some of the aircrew and aircraft resources are continuously assigned to ground alert to serve as a quick-reaction strategic force.

A rule-based system called DOSS (Decision Oriented Scheduling System) was developed by Rand to incorporate the many decision rules, SAC policies, and resource constraints that affect resource allocation of aircrews and aircraft within a typical SAC bomb wing. I use DOSS to analyze the effects, on several measures of wing performance, of alternative rules and policies for allocating aircrew and aircraft resources.

The analysis focuses on the policy issues of (a) increasing the SAC alert force, and (b) flying more training sorties, in order to estimate the extent of any additional capability that may exist within SAC bomb wings given current resources. The analysis has implications for the particular policy issues that are addressed, but its more general purpose is to illustrate the potential of DOSS for examining a broad range of resource-allocation policy issues within SAC, and possibly other Air Force Commands as well.

Annual operations and maintenance expenditures for SAC flying organizations are approximately \$1.7 billion (U.S. Congress, 1979a:71). Therefore, relatively small improvements in resource allocation could produce striking amounts of absolute dollars either

saved or turned into increased performance. In addition, many other military organizations are constructed in a fashion similar to SAC and face similar decision problems and environments. Those most closely related are the Tactical Air Command, Military Airlift Command, Air Training Command, Air Defense Command, and the fleet operations of the Navy, which account for a large part of the fiscal year 1979 national defense operations and maintenance budget of \$37 billion (U.S. Congress, 1979a:71). Methods and policies investigated in this study for improving resource-allocation policy within SAC may have direct applicability to these and other similar organizations. This investigation certainly has special policy relevance to the United States Air Force as a methodological approach to investigating and formulating corrective policy concerning resource allocation within its diverse flying organizations.

The plan for the remainder of this report is to first present the theoretical framework for the design and use of rule-based models of organizational decisionmaking. Therefore, Section II first presents a theoretical view of organizational decisionmaking that serves as a basis for the design and use of rule-based models as analytic tools.* A brief description is provided of the resource allocation process within SAC bomb wings, intended to provide further clarification and evidence supporting the theoretical basis for the design and use of DOSS as an analytic tool for SAC. Finally, a brief description of the design concept of DOSS is presented, along with examples of how DOSS

*Operationally oriented readers may wish to merely scan much of this theoretical discussion.

can be used for analysis. Section III is devoted to a more detailed description of the DOSS model, and presents evidence supporting the validity of the model. Section IV describes the design of the analysis. It outlines the issues that are addressed, the rule and policy changes that are examined, and the wing output measures that are monitored. Sections V through VII present the results of the analysis. Section VIII discusses the issue of implementing some suggested improvements in SAC wing scheduling shown possible as a result of the analysis. Finally, Section IX is devoted first to a brief summary of the major findings and policy implications of the analysis. This is followed by a discussion of some general implications of DOSS for improving resource allocation within SAC. In conclusion, some broader implications of rule-based modeling for improving organizational decisionmaking in general are discussed.

II. THEORETICAL FRAMEWORK

DECISION THEORY AND RULE-BASED MODELING

The theoretical literature on decisionmaking is vast, and involves a number of disciplines. Economics, political science, sociology, and psychology have all produced studies of the decisionmaking processes of individuals and groups of individuals influenced by organizational goals and constraints. An exhaustive review of the literature would entail a prodigious effort beyond the scope and intent of this investigation [1], but because the literature has important implications for the design of the DOSS rule-based system, upon which this research is based, a selective review is warranted [2].

Base Paradigms

In reviewing the literature on decisionmaking, one discovers at least two very different theoretical schools of thought or paradigms; the rational and the cybernetic paradigms of decision [3].

The rational paradigm is the most widespread theoretical approach in the study of decisionmaking, at both the individual and organizational level [4]. Briefly, this paradigm, and the theories and models which stem from it, are primarily based upon the assumption of "comprehensive" rational choice among a number of alternatives. "Comprehensive" rationality as defined by Simon (1957a:76) requires complete knowledge of all feasible alternatives, complete knowledge of all possible consequences resulting from each alternative, the relative likelihood of all possible consequences, and a way of

attaching a consistent set of values to all possible consequences. In this way, an optimal decision involves the choice of the alternative that leads to the greatest expected utility or value.

Furthermore, comprehensive rationality is based solely upon the notion of a single decisionmaker. The treatment of collective decisionmaking proceeds by considering the collective as an individual entity. By assuming that the collective group reaches consensus on its utility values and probabilistic estimates, the group is considered as a single entity acting as an individual.

The cybernetic paradigm offers a theoretical approach which is much less demanding on human cognitive and information-processing capabilities. The cybernetic paradigm recognizes "cognitive limits on rationality" [5] in situations involving ill-structured, complex problems. As the term "cybernetic paradigm" attempts to suggest [6], this is a theoretical approach that asserts that in facing complex decision problems, the decisionmaker relies upon common cybernetic mechanisms, such as adaptive decision rules or heuristics, in order to simplify his problem and make decisions.

Central in understanding the distinction between the rational and cybernetic paradigm is the concept of complexity in decisionmaking. Steinbruner (1974:16) describes a complex decision problem as one in which the following conditions hold:

- o Two or more values are affected by the decision;
- o There is a tradeoff relationship between the values such that a greater return to one can be obtained only at a loss to the other;
- o There is structural uncertainty: i.e., major alternatives are not clearly defined and possibly not known, and cause and effect relationships relating alternatives to their consequences are not clear and involve a great deal of uncertainty;
- o The power to make the decision is dispersed over a number of individual actors and organizational units.

Under such conditions of complexity, it has been widely observed that the assumptions stemming from the rational paradigm require such extensive calculations and information-processing capabilities on the part of the decisionmaker as to make the decisionmaking process impossibly burdensome. This argument leads to the basic assertion of the cybernetic paradigm: Decision processes must work in a much less burdensome way.

That concept was first applied in studying organizational decisionmaking by Simon (1957a), and elaborated further by March and Simon (1958) and Cyert and March (1963). Common throughout these early theoretical works was the notion of "bounded rationality" or "cognitive limits on rationality." The physical and psychological limits of man's capacity as alternative-generator, information-processor, and problem-solver constrain the decisionmaking processes of individuals and organizations. Because of these bounds, intendedly

rational action requires simplified models that extract the main features of a problem without capturing all its complexities (Allison, 1971:71).

Based on the notion of "bounded rationality," these organizational theorists have developed such concepts as standard operating procedures, satisficing, factorization, quasi-resolution of conflict, and uncertainty avoidance, which together constitute a theoretical approach to organizational decisionmaking different from that of the rational paradigm. These aspects of decisionmaking within organizations will be described in more detail later, when a brief description of the resource allocation process within SAC bomb wings is presented.

Both theoretical approaches are widely represented in the literature and have important implications for policy analysis. As pointed out in the literature, particularly by Allison (1971) and Steinbruner (1974), different approaches taken in analyzing decisionmaking within organizations may lead to vast differences in interpretations and predictions.

The rational paradigm is largely normative in nature, depicting how decisions ought to be made while largely neglecting human cognitive limitations and commonly observed decision strategies. When the nature of the analyst's task is to model decisionmaking within an organization for purposes of prediction, methodologies and assumptions stemming from the rational paradigm may be totally inappropriate in all but the most simplistic of problem areas. It may be analytically attractive to build models based on rational assumptions to predict organizational behavior. However, the assumptions employed regarding

organizational decisionmaking and human cognitive capabilities may so oversimplify the problem structure as to lead to very faulty predictions (Strauch, 1974).

Nevertheless, many well-developed management science techniques and methodologies (mathematical programming, decision analysis, game theory, etc.) based on assumptions of rational decisionmaking exist for the analyst interested in improving decisionmaking within organizations. Techniques such as linear programming, for example, have been very successful in improving decisionmaking in organizations, but only when certain conditions are met. The problem being faced must have clearly defined alternatives, clear cause and effect relationships which can be expressed quantitatively, and clearly defined objectives which can also be expressed quantitatively. The number of variables needed to be considered cannot be too large and certain parameters have to be estimatable. Lastly, the problem must be small enough that the calculations can be carried out in reasonable time and at a reasonable cost (Simon, 1960:17).

Unfortunately, many decision problems within organizations do not meet these conditions. In more complex problems where these conditions are not met, such approaches have not yielded great success. In fact, attempting to apply such techniques to complex problems generally requires such simplification of the problem structure as to make the problem being solved by these sophisticated mathematical techniques too simplistic to be of any practical value. It is one thing to hypothesize a well-defined objective function; it is quite another thing to actually obtain one from an organization, especially in the public sector, where goals are much harder to define.

Given the limitations of methodologies stemming from the rational paradigm for modeling or improving decisionmaking within organizations, this naturally raises the question of what alternative methods can the analyst turn to. What techniques and methodologies, if any, should the analyst apply in modeling decisionmaking within an organization for the purpose of prediction, when such decisionmaking clearly involves cybernetic decision processes to deal with complexity? What techniques and methodologies should the analyst apply in attempting to improve decisionmaking within an organization when the nature of the problem is so complex that methodologies stemming from the rational paradigm are inadequate?

Unfortunately, the cybernetic literature does not provide any direct answers to these questions. The literature is almost purely descriptive in nature. Although the cybernetic literature describes common decision processes found to exist in many types of organizations, it is largely inadequate when it comes to suggesting ways for improving decisionmaking in organizations. Nevertheless, a highly related field of research does have potential in this regard: the field of artificial intelligence.

Rule-Based Modeling

Works stemming from Simon (1959, 1960, 1966, 1968), Newell and Simon (1958, 1972), Reitman (1965), Feigenbaum (1977), and others, have addressed the issues of simulating human decision processes in order to predict behavior and serve as an aid in improving decisions. Rule-based modeling, as stated previously, is based on the assumption

that many types of complex human problem-solving can be broken down into an extensive set of interrelated heuristics or decision rules. Rule-based modeling of decisionmaking within organizations is therefore based on assumptions that are highly consistent with the cybernetic view of organizational decisionmaking.

The fundamental idea underlying rule-based modeling, for analysis purposes, is that in order to improve decisionmaking you first have to understand it. In this regard, rule-based modeling seeks to identify the heuristics or decision rules that decisionmakers actually use. With the aid of the computer, the rules and their interrelationships can be programmed to simulate actual decisionmaking behavior. Once a valid descriptive model that reproduces final output is developed, the model can then be used as an analytic tool to depict the effects on output that result from alternative decision rules.

Rule-based modeling has been used to develop and test the validity of descriptive models of decisionmaking in a number of organizational contexts. Cyert and March's (1963) behavioral model of the decisionmaking processes of a private-sector business firm, is perhaps the best-known model that has been used to simulate organizational decisionmaking. In this pioneering work, Cyert and March programmed into a computer the decision rules that they theorized reflected the price and output decisionmaking behavior of a retail department store. They then tested their theories by comparing the outputs of their model to actual decisions made by the firm. In addition, they also suggested how such a model could be used in an analytic fashion to examine the effects of alternative decision rules. Clarkson (1962) has applied rule-based modeling in simulating

portfolio selection by trust officers of a bank, which has long been a decision process where optimization algorithms have been applied. However, viewing trust officers as seeking optimal portfolio selections leads to very faulty predictions about their behavior when viewed empirically. Clarkson's rule-based model, however, simulated the observed behavior quite well. In the public sector, Crecine (1969) has applied a similar rule-based computer approach to build and test the validity of a descriptive model for municipal resource allocation.

Recent developments in artificial intelligence (AI) have used rule-based computer systems for the purposes of "knowledge engineering." Knowledge engineering has been defined as "bringing the principles and tools of AI research to bear on difficult applications problems requiring experts' knowledge for their solution" (Feigenbaum, 1977). Basically, these systems incorporate rules for solving complex problems that are derived from experts in their respective fields, and apply this knowledge to produce useful inferences for the user. Recent work in the fields of chemical analysis (Feigenbaum, 1971), medical diagnosis (Shortliffe, 1976), air traffic control (Wesson, 1977), analysis of international terrorism (Waterman, 1977), information management in criminal court proceedings (Buchanan, 1977), and modeling foreign policy decisionmaking (Sylvan, 1977), have shown some potential for rule-based computer systems to aid decisionmaking in a variety of contexts.

The focus of this study is on the use of rule-based modeling in the context of organizational decisionmaking, for the purposes of analyzing the effects on organizational outcomes of alternative

decision rules. The specific organizational decision problem of concern is resource allocation within bomb wings of the Strategic Air Command.

The following subsection is devoted to a brief description of the resource allocation process within SAC bomb wings. The purpose is to provide further clarification and evidence supporting the theoretical basis for the design and use of a rule-based model (DOSS) as an analytic tool for SAC. For a more exhaustive treatment of the resource allocation decision processes observed in SAC bomb wings, along with a detailed account of the organizational structure of SAC, see Berman (1974, 1975).

RESOURCE ALLOCATION WITHIN SAC BOMB WINGS

SAC's official goal for manned aircraft is to maintain mission-ready aircrews and aircraft as a credible nuclear deterrent. The Congress, through the Department of Defense, provides SAC with the resources--manpower, weapon systems, and supporting systems--to achieve this goal. This vague goal of "deterrence" has to be translated into operationally meaningful subgoals in order for SAC to function.

SAC operates decentralized bombardment wings at 19 locations within the United States, each of which conducts a flying program to maintain proficiency of aircrews and to exercise and maintain aircraft systems. Additionally, some of the aircrew and aircraft resources are continually assigned to ground alert to serve as quick-reaction strategic forces [7]. These intermediate level subgoals of alert, aircrew proficiency, and maintenance of reliable ready aircraft, are

still too vague to be operationally useful; there is no clear way of testing the extent to which these goals are realized if particular courses of action are chosen (March and Simon, 1958:155). There is no ultimate answer to how much aircrew proficiency or aircraft reliability is enough.

The wing is the operational unit of SAC. A standard wing has 14 B-52 and 14 KC-135 aircraft, which are operated by approximately 18 six-man mission-ready (MR) bomber crews and 18 four-man mission-ready tanker crews. The aircraft are prepared for flight, inspected, and repaired by nearly 1000 maintenance personnel.

Because of the complexity of jointly considering aircrew and aircraft needs, the flying and alert activities of aircrews and aircraft are determined separately by two suborganizations of the wing, using guidelines from higher echelons. The operations organization schedules crew members for sorties (training flights), ground alert, leave, and ground training; its major focus is on certain intermediate subgoals of SAC, namely aircrew proficiency and the personal needs of the aircrews. The maintenance organization schedules aircraft for sorties, inspections, alert, and repairs; its major focus is on different intermediate subgoals of SAC, namely, maintaining reliable ready aircraft and the personal needs of maintenance crews. This process of separating a complex problem into a number of parts, which are then parcelled out to various organizational subunits, is the process of factorization, commonly found in many organizations (March and Simon, 1958:192).

It is within the operations and maintenance suborganizations where operationally meaningful goals are developed that only

indirectly relate to the overall official goal of SAC: deterrence. These goals are reflected in scores of operational decision rules regarding the allocation of aircrews and aircraft for flying, maintenance, and alert activities. For example, the SAC subgoal of aircrew proficiency is reflected by numerous rules for allocating aircrews to training sorties. The SAC subgoal of maintaining reliable ready aircraft is reflected by numerous rules for allocating aircraft to flying and inspection activities [8].

These decision rules are classic examples of cybernetic processes. The development of rules, like most decisionmaking in organizations, is largely adaptive in nature, with specific decision rules or standard operating procedures developed over time to deal with a variety of recurring situations (Cyert and March, 1963:100-113).

Many of SAC's decision rules, such as sortie requirements and inspection intervals, are actual stated policy, stipulated in regulations from higher headquarters. Many other rules, however, are local wing rules developed to implement SAC policies and adhere to other subgoals within operations and maintenance. For example, a typical maintenance rule is that a B-52 aircraft cannot fly three days prior to being scheduled to go on alert. This rule reflects the maintenance subgoals of having enough time to prepare the aircraft for alert, and ensuring the aircraft will be available to meet the maintenance alert schedule. These types of rules may vary across bomb wings. For example, Berman (1975:70), in comparing the relative efficiency of different SAC bomb wings, noted that some wings spend up to five days to prepare an aircraft for alert (although it was estimated only 22 hours were actually needed).

The fact that aircraft are often not allowed to fly several days prior to going on alert denotes another process commonly found in organizations; the process of uncertainty avoidance. Comprehensively rational agents deal with alternative consequences of action by estimating probabilities of possible outcomes. People in organizations are quite reluctant to base activities on estimates of an uncertain future. They avoid the requirement of anticipating future reactions of other parts of the environment by arranging a negotiated environment. They impose plans, standard operating procedures, and uncertainty-absorbing contracts on that environment, which leads to payments of excess resources or slack (Cyert and March, 1963:118-120). Slack provides a fund of resources that can be used in an uncertain future; it acts as a buffer or hedge against uncertain outcomes (Thompson, 1967:150; Downs, 1967:138).

Keeping an aircraft on the ground long enough to guarantee (100 percent of the time) that it will be ready for its next scheduled alert, is an expenditure of excess resources (slack) in an attempt to gain certainty. Uncertainty avoidance becomes a goal in itself. The insurance that slack resources provide is not necessarily inefficient; however, such insurance does not come without a cost in terms of the availability of aircraft to fly and upon which to accomplish aircrew training.

Although operations and maintenance schedulers focus on different parts of the wing resource-allocation problem, they are required, nevertheless, to coordinate their flying activities. Since they focus on different subgoals their preferred flying schedules rarely

coincide. Through a process of negotiation and mutual adjustment a monthly schedule is planned, and each week a firm schedule is determined. It is a formalized procedure resulting in an actual flying contract signed by wing executives.

The process of negotiation and mutual adjustment to resolve conflicts, which Cyert and March (1963:117-118) refer to as the quasi-resolution of conflict, is commonly found in many organizations. Furthermore, it is postulated that excess resources (slack) will be expended to absorb the potential inconsistencies that result (Cyert and March, 1963:118). This type of behavior has also been observed by Berman (1975:72) to exist in SAC bomb wings. Slack resources introduced to resolve conflicts among operations and maintenance were seen as causing losses in terms of the availability of aircraft to fly and lowered training productivity.

As expressed above, the resource-allocation decision problem within SAC bomb wings is highly complex. It is a problem influenced by several decisionmakers and organizational units, and involves numerous organizational subgoals and constraints. Furthermore, there is an inherent tradeoff relationship between numerous subgoals. The longer an aircraft is on the ground, either on alert or in an inspection, the less it is available for flying and the accomplishment of training sorties for aircrews. Moreover, there is structural uncertainty; i.e., the cause-and-effect relationships between alternative schedules and wing performance are not clear and involve a great deal of uncertainty.

The complexity of resource allocation within SAC bomb wings leads to an additional cybernetic decision process; the process of

"satisficing." Simon's (1957a) famous satisficing model asserts that due to "cognitive limits on rationality" optimization is replaced by satisficing. In choosing, human beings do not consider all alternatives and pick the action with the best consequences. Instead, they find a course of action that is good enough; that satisfies. As mentioned earlier, a great many rules are developed in wing scheduling that stipulate what is meant by "satisfactory" achievement of a goal. Schedulers, for the most part, adhere to these rules and seek schedules that satisfy them. Because of the lack of clear cause and effect relationships, the enormous numbers of rules, and the vague nature of the goals, it is difficult enough to develop even one "satisfactory" schedule, much less search for "better" ones.

In sum, because of the complex nature of the problem, a number of simplifying cybernetic decision processes are evident in resource allocation within SAC bomb wings, such as standard operating procedures, satisficing, factorization, quasi-resolution of conflict, and uncertainty avoidance. A large number of decision rules are developed by schedulers, both in operations and in maintenance, that embody the above decision processes. These rules focus independently on operations and maintenance subgoals, and are developed with little awareness of their highly interdependent effects on the wing as a whole. Possible inefficiencies with respect to overall organizational goals may result from such processes, such as excess slack resources developed to avoid uncertainty and resolve interorganizational conflict.

The aim of this study is to view the numerous rules and policies affecting resource allocation within SAC wings as a system, which

affects wing performance. This work begins where earlier work by Berman (1975) leaves off. Identifying the types of decision processes involved in SAC wing scheduling, and the kinds of inefficiencies that may result, is only a first step to improving resource allocation. To seek improved rules and policies one needs a mechanism by which to systematically analyze the effects of alternative rules and policies. To this end, I use a rule-based model, DOSS, to first incorporate the various rules and policies prevalent in SAC wing scheduling in both operations and maintenance. I then use DOSS in an analytic fashion to examine the effects on wing outcomes of alternative rules and policies.

DESIGN CONCEPT OF DOSS

The resource allocation problem of a SAC bomb wing, as previously described, is a highly complex problem. It is a problem which clearly fits into that class of complex problems not lending themselves well to established optimization techniques for solution. Due to the enormous scale of the problem (numbers of resources, constraints, and objectives), its dynamic nature (continually changing resources, constraints, objectives), and the difficulty of ever formulating mathematically the relationships of resources, constraints, and objectives, optimization techniques are largely inappropriate. For optimization techniques to be applicable in this area, combinatorial factors alone would demand such oversimplification of the problem as to deprive the results of any practical value.

Because of the inapplicability of optimization techniques, and because of the predominance of cybernetic decision processes (many of

which can be stated as simple decision rules), a rule-based computer system, DOSS, was developed to model the resource allocation process within SAC bomb wings [9].

DOSS operates on a series of rules that reflect actual scheduling rules and SAC policies. These rules are communicated to the computer system in flexible, English-like statements. Furthermore, DOSS allows the analyst to change rules easily, and to rapidly see the consequences of rule changes in the form of operationally useful performance measures, such as total sorties that result.

DOSS was designed to accommodate both maintenance and operations rules. The following are examples of actual DOSS rules that maintenance and operations schedulers adhere to, in roughly the same form as would be input into DOSS.

o Constraints

- The activity fly cannot be assigned to a member of the set B52 if DAYS.UNTIL.ALERT are less than 3.0;
- The activity fly cannot be assigned to a member of the set B52 if FLY.HOURS.SINCE.PHASE. are greater than 100.0;

o Preferences

- The preferred member of the set aircrew for the assignment of fly has the highest SORTIES.REMAIN.;

The first rule depicted above is a constraint that prohibits an aircraft that is scheduled for alert from flying less than 3.0 days

prior to the alert. The primary purpose of such a rule is to allow enough time for maintenance to prepare the aircraft for alert. However, as discussed earlier, this rule also reflects a degree of uncertainty avoidance. With DOSS, the analyst can easily vary this rule to see how any changes will affect a number of performance measures. For example, the analyst can change the 3.0 parameter to 1.0 and thereby increase aircraft availability to fly. Then, by generating a schedule using the new rule, and holding the numerous other rules constant, he can measure the marginal effect of this increased availability on such measures as total sorties that can now be scheduled.

The second decision rule depicted above is also an example of a constraint, but in this case it is a stated SAC policy. The policy is that a B-52 aircraft cannot fly if it has flown more than 100.0 hours since its last phase inspection [10]. With DOSS, the 100.0 parameter can easily be changed to 150.0 or 200.0 to show the effects of increasing aircraft availability (as a result of lengthening this inspection interval) on a number of operationally useful performance measures.

It should be stressed that DOSS provides no answers to how aircraft reliability may be affected over time due to increased inspection intervals. It provides only a methodology for estimating the possible benefits of such a policy change; for example, the additional sorties that can be scheduled as a result of increased aircraft availability. Other means of analysis would be needed to completely analyze the worth of any such policy option, if indeed the benefits would seem to warrant further investigation.

The third rule depicted above is an example of a preference rule. This is a type of rule for choosing particular resources for an activity that do not violate any constraints. The rule stated above expresses the preference of operations schedulers to choose an available aircrew to fly that has the highest number of sorties remaining until its sortie requirements are met. This reflects the goal of scheduling an equal distribution of sorties for all aircrews.

Taken together, constraint and preference rules, of the type discussed above, are the major rules incorporated in DOSS for allocating resources to various activities.

Having discussed the theoretical basis for the design of DOSS, and a few examples of how it can be used as an analytic tool, a more detailed description of the model will be presented next. This will be followed by a description of the analysis that is performed with DOSS, and several sections reporting the major results.

NOTES--SECTION II

1. A number of reviews are already available; see Becker and McClintock (1967), Edwards (1961), Johnson (1968).
2. For a more complete discussion of decision theory and its relevance for the design of computer-based decision support systems, see Keen and Morton (1978).
3. See Steinbruner (1974). I use the term "rational" paradigm instead of the term "analytic" paradigm as introduced by Steinbruner, to denote the basic concept of rational choice underlying the paradigm, and to avoid some of the difficulties with the word "analytic." As I argue in this study, other views of decisionmaking hold promise for "analytic" purposes as well.
4. For a rigorous treatment of the concepts underlying this paradigm, see von Neuman and Morgenstern (1954), Luce and Raiffa (1957), Luce and Suppes (1965).

5. "Cognitive limits on rationality," a term first phrased by March and Simon (1958), refers to the limit of human cognition which prohibits solving a complex decision problem in a comprehensively rational way.
6. The term cybernetic refers to biological organisms and mechanical mechanisms that are highly adaptive in nature, and rely on rapid information feedback to allow for corrective action (e.g., a house thermostat). See Steinbruner (1974:48-56) for a detailed discussion of this concept.
7. Currently, approximately 30 percent of the bomber force is on continuous ground alert (U.S. Congress, 1979b:397).
8. Berman (1975:19) has identified approximately 60 operational goals within operations and maintenance organizations of SAC bomb wings. Most of these are stipulated as simple decision rules to achieve a certain measurable quantity (such as sorties), or not to exceed a certain measurable quantity (such as days since last inspection). Furthermore, many other rules exist for scheduling day-to-day activities related to these goals. For an illustrative example of the many scheduling rules involved in the maintenance area within SAC, as well as in other Air Force Commands, see Leger and Peacock (1975).
9. The principal developer of this computer system has been Dr. R. J. Hillestad of The Rand Corporation. The rule-based models of SAC wing scheduling used in this study were developed by this author, with the assistance of numerous SAC schedulers, employing the DOSS computer system.
10. A phase inspection is a scheduled inspection of the major systems of the aircraft. Its duration is approximately one week.

III. MODEL DESCRIPTION AND VALIDATION

MODEL DESCRIPTION

The purpose of the DOSS rule-based model of SAC wing scheduling is to examine the effects of alternative scheduling rules and higher-headquarters SAC (Hq SAC) policies on many measures of wing performance. Therefore, DOSS was developed to incorporate, as closely as possible, the major scheduling rules, Hq SAC policies, and resource constraints that operations and maintenance schedulers adhere to. The purpose was to establish a baseline model from which comparisons of alternatives could be made.

This work was conducted within the context of the then ongoing B-52 Aircrew Continuation Training Test Program, which took place at selected SAC bomb wings from September 1978 through February 1979, and has subsequently been implemented on a command-wide basis [1].

In developing a set of DOSS rules that accurately reflect aircrew and aircraft scheduling, several visits were made to two SAC bomb wings, who were the original participants in the test program. For the purposes of this study, I refer to the two test wings as Wing A and Wing B. During these visits both operations and maintenance schedulers described in great detail the rules and policies they adhered to in order to implement the test program. I used this information to develop sets of DOSS rules, called processes, that actually generate schedules of aircrews and aircraft for various types of training sorties. Since the objective was to develop a representative or baseline model of such scheduling processes for the purpose of examining alternatives, only one of the two wings visited

was modeled in great detail. I used the many rules derived from Wing A in developing what I refer to as the DOSS baseline model.

What follows is a brief description of the major factors (decision rules, higher-headquarters policies, and resource constraints) that influence scheduling of B-52 aircrews and aircraft for various types of training sorties, which I incorporated as DOSS rules in the baseline model. My intent is to clarify these major factors, and to illustrate the level of detail incorporated by the DOSS baseline model.

Aircrew Scheduling

A primary function of operations schedulers is to schedule mission-ready (MR) aircrews to training sorties. Hq SAC defines three distinct classes of sorties required to be flown: a profile training sortie [2], an event training sortie, and a pilot proficiency sortie. Each sortie is distinguished by the types of training events it contains, its approximate duration, and the number of assigned crew members required to participate [3]. The typical profile sortie, of approximately seven hours in duration, requires a bomber and tanker meeting for air-refueling practice, and the bomber crew practicing low-level navigation and low-level bombing exercises against a radar bomb-scoring (RBS) site. The typical event sortie, of approximately five hours in duration, contains low-level navigation and bombing, but not necessarily air-refueling. The pilot proficiency sortie (pilot pro), of approximately three hours in duration, is designed to maintain the proficiency of MR pilots in strictly pilot

responsibilities, such as training in instrument approaches and emergency procedures.

Hq SAC establishes four three-month training periods, or quarters, per year. In each quarter every MR crew is required to fly a certain number of sorties. In the case of MR crews available for the entire three-month period the requirements are: four profile sorties, seven event sorties, and one pilot pro sortie, for a total of at least twelve sorties per MR crew per quarter [4].

A primary objective of operations schedulers is to satisfy the above equal sortie requirements per MR crew, without violating Hq SAC policies and resource constraints (e.g., available aircrews, flying hours, and RBS site availability). Furthermore, operations schedulers must also schedule aircrews to ground alert, leave, and various ground training activities which compete for the aircrew resource.

In addition to the three classes of sorties defined by Hq SAC, wing schedulers further distinguish different types of sorties based on approximate takeoff time. The takeoff time of a sortie is important because it affects the availability of crews to fly the sortie, as well as tanker and RBS site availability, which directly affects the training content of the sortie. The typical five-hour event sortie is therefore further classified into two distinct types: (1) the "Cold Seat Swap" (CSS) sortie (usually an afternoon takeoff around 1400), and (2) the non-CSS event sortie (usually a morning takeoff around 0800). The CSS sortie, although it differs in training content from that of the non-CSS event sortie, primarily affects maintenance scheduling, and will be discussed in more detail in this context later. Schedulers further distinguish profile sorties as

either night profiles (takeoff around 2200) or day profiles (takeoff around 0800), again because of differences in aircrew availability and training content.

The scheduling process takes place in three distinct phases in which a quarterly, monthly, and finally a firm weekly schedule are produced. Prior to the beginning of each three-month training period, the operations scheduler produces a quarterly schedule. In the quarterly schedule he first assigns aircrews to their alert duty and their leave. Having accomplished this, he then attempts to schedule MR crews to sorties of different types taking into account any other known activity (e.g., TDYs--temporary duty away from the bombing squadron).

In scheduling sorties to MR crews, the scheduler follows what I refer to as a horizontal scheduling strategy. Basically, the scheduler attempts to fill in a matrix with scheduled flying activity where the columns of the matrix are the days of the quarterly period and the rows represent distinct crews. The scheduler first schedules flying activity by considering one crew at a time. He moves horizontally across a particular crew's row until he schedules all required training sorties for that crew. In so doing, the scheduler follows a set of rules that allow each MR crew the required time between sorties for preflight crew rest and mission planning [5]. Also, he follows additional rules in scheduling standardization/evaluation crews (stand/eval), in order to allow sufficient availability for such crews to perform evaluations of other crews in flight [6].

The difficulty with this scheduling strategy arises near, or at the end, of having scheduled all the crews, or rows of the matrix. The difficulty is due to the fact that the operations scheduler is constrained to schedule a certain number of sorties, agreed upon in advance through negotiation with maintenance, for each day (or vertical column) of the period. This agreement is referred to as the wing's sortie flow, and reflects constraints posed by the capability of maintenance to generate sorties [7]. At the end of a horizontally built schedule, the number of sorties scheduled for a particular day (or vertical column) will invariably not match with the number of sorties contracted for with maintenance that are planned to be generated on that day. Because of numerous mismatches between what is scheduled and required on a daily basis, the scheduler then proceeds to make marginal changes in the schedule, wherever possible, in order to match the number of sorties scheduled each day with the requirement. In the process of so doing, however, the number of sorties previously scheduled to crews will change, and therefore may no longer represent a "satisfactory" distribution of training sorties. As mentioned earlier, schedulers attempt to distribute to each crew a minimum of 12 total sorties of different types, in accordance with Hq SAC policy. Therefore, the operations scheduler usually makes several iterations, sometimes taking several tedious hours, until he produces a "satisfactory" quarterly schedule.

Prior to each month the operations scheduler also produces a monthly schedule. This schedule reflects any new conditions not known at the time when the quarterly schedule was planned. Such items as additional sorties to make up for previous sortie losses, and higher-

headquarters directed missions (HHDMs) [8] for the month are factored into the schedule at this time.

Finally, prior to the beginning of each week (when all activities, such as alerts, HHDMs, TDYs, and leaves, are relatively fixed), the scheduler makes a firm assignment of crews and crew members [9] to flights in the weekly schedule. In addition, the weekly schedule includes a great deal of specific information about the sorties scheduled (e.g., type of sortie, training items to be accomplished, specific takeoff time, route to be flown, RBS site or sites, and more).

The DOSS baseline process for aircrew scheduling incorporates in great detail most of the major factors discussed above, and many others as well. The model essentially takes as given the known activities and resource constraints facing operations schedulers prior to producing weekly schedules. It then attempts to allocate sorties of the different types discussed above (i.e., non-CSS event sorties, CSS sorties, day and night profile sorties, and pilot pro sorties) to MR crews in order to fulfill overall sortie requirements, as well as HHDM responsibilities, for a three-month time period.

More specifically, the model takes as given the alert, leave, TDY, and ground training activities of all MR crews (along with any 4x sorties scheduled to accomplish staff training [10]), based on the weekly schedules of Wing A for the quarterly period September through November 1978. In addition to these constraints on aircrew availability for flying, the DOSS baseline model also incorporates many other rules stipulating the wing's typical sortie flow (which reflects constraints posed by maintenance), along with rules regarding

how this "flow" varies on days when HHDMs are required. The model then attempts to allocate sorties to MR crews on a daily basis for the period September through November 1978, comprising a total of thirteen complete weekly schedules.

In scheduling sorties to MR crews, however, the DOSS baseline model by design follows what I refer to as a vertical scheduling strategy, in contrast to the horizontal strategy that schedulers typically follow. The vertical strategy essentially treats each day, or vertical column of the scheduling matrix, independently. DOSS schedules the particular daily requirement of sorties contracted for with maintenance to those available crews which have the most "need" for the sorties. A crew is considered available if it meets a host of constraint rules, such as whether sufficient time has elapsed since the crew's last sortie for crew rest and mission planning. A crew is considered to have the greatest "need" for a particular type sortie on the basis of a number of preference rules. For example, the available crew with the greatest "need" for a type A sortie will be that crew with the highest number of type A sorties remaining until it fulfills its type A sortie requirement.

Using the vertical strategy, a routine set of rules are systematically followed for each day of the scheduling period. Therefore, no one day is any more difficult to schedule than any other. Furthermore, at the end of the scheduling period all daily sortie requirements should be fulfilled, and the distribution of sorties to MR crews should be highly correlated with sortie requirements. The distribution of sorties should be correlated with sortie requirements, because over time if a crew falls behind in terms

of sorties flown relative to other crews (because of high amounts of alert or leave), then that crew becomes the most preferable to fly until it catches up.

I programmed the DOSS to follow a vertical scheduling strategy in order to examine the effects such a strategy would have on possibly improving sortie distributions. I will address in detail the analysis of sortie distributions produced with DOSS, as compared with those that wing schedulers produced, in later sections.

The DOSS baseline model does not schedule individual crew members to sorties because this level of detail, encompassing day-to-day scheduling functions, was not necessary for the type of issues I address with the model. Furthermore, DOSS attempts no rescheduling of crews to make up for any sortie cancellations (which, for the most part, was also the case at both wings which were studied). Nevertheless, the extent of differences between what is scheduled on a weekly basis and eventually accomplished (the differences are not great) may have implications for the analysis discussed later, and will be dealt with separately at that time.

Aircraft Scheduling

A primary function of maintenance schedulers is to schedule B-52 aircraft to training sorties. Maintenance schedulers, like operations schedulers, schedule in three distinct phases in which a quarterly, monthly, and finally a firm weekly flying schedule are produced. The other activities that maintenance is responsible for, which are usually firmly set prior to the development of the weekly flying schedule, are the alert, phase inspection, programmed depot

maintenance (PDM) [11], time compliance technical orders (TCTO) [12], and weapons load training (WLT) [13] activities of aircraft. All of these activities compete for the aircraft resource.

Maintenance schedulers also face a high degree of uncertainty as to when aircraft components will fail and how long repairs will take. Thus, they adhere to many rules allowing adequate ground time for an aircraft between sorties, in order to clear any "inflight discrepancies" so that it will be ready for its next sortie.

Maintenance schedulers classify different types of sorties based on duration and approximate takeoff time, and thus take into consideration the five types of sorties mentioned earlier (non-CSS event sorties, CSS sorties, night and day profile sorties, and pilot pro sorties). As discussed earlier, a sortie flow, already agreed upon with operations, stipulates how many sorties of different types maintenance is responsible to schedule each day.

In scheduling different types of sorties, maintenance schedulers adhere to many rules that allow sufficient time between sorties for accomplishing a preflight and basic postflight (BPO) inspection. These inspections are required before and after every regular (or non-turnaround) sortie. A turnaround sortie is a sortie where the aircraft flies only a short while after it has already flown a previous sortie, and therefore does not require a BPO following the previous sortie or a preflight inspection prior to the new (or turnaround) sortie. However, prior to most turnaround sorties sufficient time is allowed in order to perform a thru-flight inspection [14]. Furthermore, maintenance schedulers adhere to other rules that reflect their preferences in certain situations for

scheduling non-turnaround sorties (requiring a preflight and BPO inspection) or turnaround sorties (usually requiring a thru-flight inspection) [15].

The CSS sortie is a unique type of sortie. It essentially is a very quick turnaround sortie. This is a sortie that typically uses one of two B-52 aircraft that have just landed from a previous non-CSS event sortie, with no more than one and one-half hours ground time in between. During the maximum of one and one-half hours, only very minor maintenance is performed. The idea behind the CSS sortie (which was one of the innovations of the B-52 Aircrew Continuation Training Test Program) was to generate more sorties with current resources, without increasing maintenance workloads.

The DOSS baseline process for aircraft scheduling takes as given the known activities and resource constraints facing maintenance schedulers prior to producing weekly flying schedules. It then attempts to allocate sorties, of the different types discussed above, to B-52 aircraft in order to fulfill weekly sortie commitments made to operations, as well as HHDM responsibilities.

More specifically, the model takes as given the alert, phase inspection, PDM, TCTO modifications, WLT, and any other planned activities of all B-52 aircraft of the wing, based on the weekly schedules of Wing A for the quarterly period September through November 1978. In addition to these constraints on aircraft availability to fly, the DOSS baseline model incorporates many other rules stipulating the wing's typical sortie flow, along with rules regarding how this "flow" varies on days when HHDMs are required. The model also incorporates many rules in order to meet the stated

preferences of maintenance schedulers regarding the number of aircraft to fly each week, and how often they are to be used for turnaround sorties. The model then attempts to allocate sorties to B-52 aircraft on a daily basis for the period September through November 1978, comprising a total of thirteen complete weekly schedules.

The DOSS baseline model makes no attempt at rescheduling aircraft in order to deal with any sortie cancellations and substitutions due to aircraft malfunctions. Nevertheless, the extent of differences between what is scheduled on a weekly basis and eventually accomplished (the differences are not great) may have implications for the analysis discussed later, and will be dealt with separately at that time.

MODEL VALIDATION

Whether a model of any process is valid or not depends upon the purposes for which it is used. As mentioned above, the DOSS baseline model does not incorporate rules for individual crew member scheduling, nor does it incorporate rules for rescheduling crews, crew members, and aircraft. Therefore, it does not reflect all the actual functions of wing schedulers which are involved in finalizing any schedule and seeing it actually flown; nor is its purpose in this investigation to act as such a scheduling device. Rather, its purpose in this investigation is to incorporate the major scheduling decision rules, Hq SAC policies, and resource constraints that primarily affect wing output in terms of total sorties and alert, in order to provide an analysis tool to examine alternatives. For example, the analysis focuses on the extent of any additional capability within a typical

SAC bomb wing, given current resources, to (a) fly more training sorties, and to (b) increase the alert force. Since the scheduling of individual crew members to particular sorties does not significantly affect the levels of total sorties and alert that are feasible, I did not incorporate crew member scheduling in the model. Also, the extent of any differences between what is scheduled on a weekly basis and eventually accomplished (again, the differences are not great) is dealt with separately.

In order to validate that the DOSS baseline model adequately reflects the many interactions of decision rules, Hq SAC policies, and resource constraints that largely determine what is scheduled on a weekly basis, I compared the outputs of the DOSS baseline model with the outputs of actual weekly schedules. I compared what the DOSS baseline model scheduled for the period September through November 1978, with what actually was scheduled on a weekly basis at Wing A during the same time period and under identical resource constraints.

Table 1 depicts the mix of sorties that the DOSS baseline model scheduled to all MR crews, and the mix of sorties that wing schedulers at Wing A scheduled on a weekly basis. It essentially is a comparison of what DOSS scheduled, using the same rules that schedulers at Wing A said they attempted to adhere to, with what actually was scheduled at Wing A. As Table 1 depicts, the two mixes of sorties scheduled were very similar. Only on 4 out of 67 flying days (which takes into account holidays) during the period does the DOSS baseline model not schedule the same number and type of sorties that were actually scheduled at the wing.

Table 1
SORTIES SCHEDULED
(September-November 1978)

	^a Wing	DOSS Baseline
Total sorties	239	241
^b Flying hours	1377	1383
Non-CSS event sorties	76	79
CSS sorties	46	47
Night profile sorties	57	58
Day profile sorties	47	44
^c Separate pilot pros	13	13

^a
SOURCE: Weekly operations schedules of Wing A.

^b
Flying hours are approximated using the averages of 5 hours per non-CSS event sortie, 5 hours per CSS sortie, 7 hours per profile sortie, and 3 hours for a pilot pro.

^c
Some pilot pros are dual logged on day profile sorties. With DOSS, 6 were scheduled this way for a total of 19 pilot pros; at the wing 10 were scheduled this way for a total of 23 pilot pros.

Table 2 compares maintenance activities directly related to flying that the DOSS baseline model scheduled, with those activities scheduled on a weekly basis at Wing A. As with the case of sorties scheduled, we see a very close similarity. The DOSS output, as mentioned earlier, reflects the stated preferences of maintenance schedulers for choosing the number of aircraft to fly each week (either 7 or 8 aircraft are preferred depending on special circumstances), and their preferences for choosing how often those

aircraft are to be used for turnaround sorties. The very slight differences shown in Table 2 result from schedulers not adhering to their own stated rules, either because of the complexity of doing so, or because they were adhering to additional rules not accounted for in the model to deal with unusual circumstances.

Table 2
SCHEDULED MAINTENANCE ACTIVITIES RELATED TO FLYING
(September-November 1978)

	(a) Wing	DOSS Baseline
Average number of aircraft scheduled to fly per week(b)	7.69	7.46
Preflight inspections	110	108
Turnarounds	83	86
BPOs	99	96

a
SOURCE: Weekly maintenance flying schedules of Wing A.

b
Excludes flying activity on 9/01; 13 complete weeks.

On the basis of the very close similarity shown in Tables 1 and 2, on measures related to total sorties scheduled, the DOSS baseline model appears valid for the purposes for which I use it in this study. Again, the purpose of the baseline model was to adequately reflect the many scheduling rules, Hq SAC policies, and resource constraints that affect wing output in terms of sorties and alert. In so doing, the baseline model provides a useful analytic tool to examine the effects on wing output resulting from alternative rules and Hq SAC policies.

Some significant differences do exist, however, between what the DOSS baseline model scheduled and what was scheduled at the wing, with regard to the distribution of sorties to MR crews. Some differences were to be expected, however, because the DOSS baseline model and wing schedulers use different scheduling strategies in distributing sorties to MR crews. As discussed earlier, the DOSS uses a vertical strategy, and wing schedulers follow a horizontal scheduling strategy. The differences in sortie distributions are important for some issues addressed later regarding alternative policies for distributing sorties to crews, and will be discussed in detail at that time.

NOTES--SECTION III

1. The SAC B-52 Aircrew Continuation Training Test Program was a test program where a number of significant policy changes were made in an attempt to increase overall wing effectiveness at no significant increase in cost. The policy changes of this program are discussed in more detail in Section IV. For a detailed description of the test program see "Strategic Air Command Program Plan 78-10."
2. There are actually two types of profile sorties: an EWO Profile Training Sortie (EPTS) and a regular profile training sortie (PTS). The EPTS has slightly higher training content, but for the purpose of this investigation this distinction is not significant. (See SAC Program Plan 78-10:178.)
3. All sorties are meant to be integral crew sorties, where all members of the same crew participate. Substitutions are permitted in accordance with specific rules, which are different for different types of sorties. (See SACM 51-52, Vol. IV:1978.)
4. In addition to MR crew sortie requirements, there are also requirements for many specific training items for different crew positions. (See SAC Program Plan 78-10:1978.) For example, 18 BOIs (Low Altitude Multiple Release Bomb Runs) are required for each MR radar navigator during the three-month training period. Because of the typical training

content of the various types of sorties, if a crew gets its required share of sorties usually each crew member will fulfill his individual requirements. However, occasionally a crew member may be ill or on emergency leave for some extended period of time, and his requirements then have to be dealt with on an individual basis.

5. Prior to each sortie, a certain amount of time is devoted to planning the details of the sortie, such as the training items to be accomplished, the route to be flown, the area where refueling is to take place, and more.
6. Usually three crews per bomb wing are designated stand/eval crews which evaluate other MR crews in flight (SACM 60-4 check), and also must meet their own MR training requirements.
7. For example, the typical sortie flow at Wing A was five sorties of specified types Monday through Wednesday, four sorties on Thursday, and two on Friday for a total of 21 sorties. This typical flow changes when special circumstances arise, in accordance with well-defined scheduler's decision rules.
8. HHDMs are required sorties demanded from Hq SAC. These sorties involve wings in various types of SAC and Air Force exercises. This class of sortie varies widely in terms of training content. When they are required (usually less than five times per month) they disrupt the typical weekly sortie flow.
9. As previously mentioned, crews normally are planned to fly as an integral unit; however, substitutions of individual crew members are allowed to some extent. Usually weekly crew assignments are stable, yet often individual crew members will become DNIF (Duty Not to Include Flying) because of illness or injury, requiring individual crew member substitutions.
10. 4x sorties are sorties where less than a specified number of personnel on board the flight belong to the same crew. When these sorties are scheduled on a weekly basis they usually are for providing wing staff members with their flying requirements. These requirements are met either on 4x sorties or by substituting staff members on regular MR crew sorties on a limited basis.
11. Every aircraft in the fleet is scheduled to depots at specific intervals and may be gone one or more months.

12. TCTOs refer to those times when modifications are made to specific aircraft systems. The wing schedules a specified period of time to complete the modifications on each aircraft.
13. Usually one B-52 aircraft is held out of the flying schedule each week to be used for training maintenance personnel in a variety of procedures.
14. Prior to the beginning of a non-turnaround sortie a pre-flight inspection of approximately six hours is required, and upon landing a BPO inspection of approximately eight hours is required. A turnaround sortie is one where the aircraft flown does not undergo a preflight inspection prior to the following flight. An aircraft qualifies for a turnaround if less than 72 hours has elapsed since its last preflight inspection (a Hq SAC policy). Therefore, this 72-hour period is that period of time in which an aircraft can be flown, or "turned," from a previous sortie without requiring a BPO following the previous sortie and a preflight prior to the new sortie. After 72 hours the aircraft must then undergo a BPO and a preflight inspection before it can fly again. Prior to a turnaround sortie, a thru-flight inspection of approximately four to six hours is usually performed, although not necessarily required.
15. Maintenance schedulers further classify turnaround sorties into those that involve more than 24 hours ground time between flights (normal turnarounds), and those that involve less than 24 hours but more than 12 hours ground time (quick turnarounds). As a general rule, maintenance schedulers prefer normal turnarounds because they allow more ground time between sorties.

IV. DESCRIPTION OF THE ANALYSIS

SCOPE OF THE ANALYSIS

This study examines the potential of a rule-based model, DOSS, as an analytic tool. To this end, DOSS is used to investigate the effects of alternative decision rules and Hq SAC policies on resource allocation at a typical SAC bomb wing. The focus of the analysis is on the capability of SAC bomb wings, given current resources, to (a) fly more training sorties, and to (b) increase the SAC alert force. The analysis has implications for the particular alternatives which are examined, but its more general aim is to illustrate the potential of DOSS for examining a broad range of other resource allocation policy issues within SAC, and possibly within other Air Force Commands as well.

The DOSS model development and analysis upon which this study is based was performed within the context of the then ongoing SAC B-52 Aircrew Continuation Training Test Program. As noted earlier, this was a test program where a number of significant policy changes were made in an attempt to increase overall wing effectiveness at no significant increase in cost. At the time of this writing, the test program had been implemented at most bomb wings throughout the command. Briefly, the new flying program is meant to increase total sorties over the old flying program by approximately 25 percent. This increase in sorties is attempted, however, without increasing total flying hours expended, and without increasing any required maintenance activities. This apparent increase in sorties at no significant increase in cost is accomplished by decreasing the flying-hour

duration of many sorties without, however, significantly decreasing the individual training items accomplished per sortie.

One set of policies whose effects on wing performance I examined with DOSS, but which I will not discuss in detail here, were the policy changes adopted under the B-52 Aircrew Continuation Training Test Program. Specifically, these policies were designed to increase aircrew and aircraft availability to fly by (1) increasing the valid preflight period from 48 hours to 72 hours, (2) increasing the maximum period an aircraft could be on alert from 42 to 84 days, and (3) decreasing the amount of time crews spend in mission planning. In addition, I examined specific differences in wing-level policies and decision rules adopted by Wing A and Wing B (the two original participants of the test program). A more detailed discussion of these examinations is reported in earlier work (Fallon, 1979).

The focus of this study is on a broader set of policy variables, chosen primarily to estimate the extent of any additional capability to achieve objectives similar to that of the test program, namely, increasing output at no significant increase in cost. For example, the analysis focuses on the extent to which it is possible, given current resources, to schedule even more sorties than are scheduled under the test program, without increasing flying-hour allocations or required maintenance activity. Furthermore, since the number of aircrews and aircraft on alert is another important measure of wing output, the analysis also focuses on the capability of SAC wings, given current resources, to increase the alert force without decreasing sortie levels. In addition, I examine the feasibility of an alternative policy for allocating sorties to MR crews, designed to

improve overall training effectiveness given a fixed number of sorties. The alternative policy is to provide different numbers of sorties to MR crews based on a measure of their relative proficiency. I refer to this alternative policy as differential training, which is in contrast to the current policy of giving each crew an equal number of sorties regardless of proficiency.

More specifically, I attempt to schedule increased sortie and alert levels with DOSS, both with and without specific rule and policy changes to the DOSS baseline model. The intent is to estimate the feasibility and consequences (on many measures of wing outcomes) of increasing sortie and alert output, both with and without the rule and policy changes. The specific changes made to the baseline model, which will be discussed shortly, were chosen primarily to increase aircrew and aircraft availability to fly.

As we observed earlier, and has been discussed in detail by Berman (1975), wing scheduling rules and policies reflect a certain amount of slack resources used to avoid uncertainty or to reduce conflicts within the organization. As mentioned earlier, a certain amount of slack may not be inefficient, but nevertheless it does represent a cost to the wing, usually in terms of aircrew and aircraft availability to fly. Some of the rule and policy changes made to the baseline model that increase aircrew and aircraft availability reflect different levels of slack resources. Therefore, in some cases I use DOSS to depict the increased output that is possible with less slack resources, or in other words the opportunity costs of the current levels of slack.

Table 3 lists the specific policy variables of the DOSS baseline model that are changed either individually or in combination, holding all other factors affecting wing output constant. These selected variables are by no means an exhaustive list of factors that determine wing output. The specific nature of each of these variables, how they affect aircrew and aircraft availability, and their effects on many measures of wing output will be discussed in detail throughout the next several sections.

The effects of the rule and policy changes made to the baseline model are examined at each of three different sortie and alert levels: (1) the current or baseline sortie and alert levels, (2) the baseline level of sorties but with increased alert, and (3) increased sortie and alert levels. The intent of such an experimental design is to estimate the feasibility and the effects of achieving (a) increased sortie levels, (b) increased alert levels, and (c) significant differential training, both with and without the changes to the DOSS baseline model.

Table 3
SELECTED POLICY VARIABLES

- o Flying before alert and after alert
- o Use of CSS sorties
- o Alternative mixes of sorties while holding flying hours constant
- o Length of mission planning
- o Differential training
- o Use of turnarounds
- o Flying hours between phase inspections
- o Length of the valid preflight period

Table 4 depicts the wing output measures that I examine for each alternative schedule generated with DOSS under different sets of conditions. These measures are classified in terms of training output, alert output, maintenance output, and morale considerations. This classification does not exhaust the list of relevant wing outputs. These measures were chosen to reflect the major effects likely to result from changes made to the DOSS baseline model, such as the increased number of aircrews and aircraft on alert that are possible due to policy changes increasing aircrew and aircraft availability.

The wing output measures depicted in Table 4 are related to overall wing effectiveness in meeting intermediate-level SAC goals, such as aircrew proficiency, alert readiness, and maintaining reliable ready aircraft. However, each of the individual output measures, such as total sorties scheduled per crew, is only one of many determinants

of wing effectiveness. For example, wing effectiveness can be thought of as some function of aircrew proficiency, alert readiness, extent of reliable ready aircraft, and morale. Furthermore, each of these intermediate SAC goals is a function of many measurable and some unmeasurable quantities. For example, aircrew proficiency is a function of many items of training output such as total sorties per crew, distribution of sorties to crews, individual training items accomplished, extent of crew coordination, diversity of sorties, and many more. It should be emphasized that because of the subjective nature of overall wing goals, these functional relationships are largely judgmental in nature. For example, the net effect on overall wing effectiveness of increasing total training sorties at the expense of decreasing time spent by aircraft in maintenance is highly judgmental.

Because of the highly complex and judgmental nature of evaluating overall wing effectiveness, I make no attempt to integrate the various output measures that result from alternative schedules into an overall measure of wing effectiveness. Instead, each output measure that is affected by a particular schedule will be presented along with all other output measures, in order to provide a balanced view of the tradeoffs involved. However, assumptions will be made, whenever appropriate, as to those measures that are likely to increase or decrease effectiveness related to specific goals. For example, increasing training sorties, everything else being equal, will increase aircrew proficiency and thereby wing effectiveness.

Table 4

WING OUTPUT MEASURES

Training Output	Maintenance Output	Alert Output	Morale Considerations
Total sorties	Preflight inspections	Number of air- crews and aircraft on alert	Average crew alert rate
Mix of different types of sorties	BPO inspections		Back-to-back alerts
Distribution of sorties to MR crews	Turnarounds		Distribution of alert
Total training items	Phase inspections		Flights before alert
Distribution of various train- ing items to MR crew members	Aircraft hours spent in sched- uled inspections		Flights after alert
Diversity of RBS sites to MR crews	Average hours ground time between sorties		Back-to-back sorties
Total number of night sorties	Average number of aircraft scheduled to fly per week		Average duty workweek per MR crew
Average number of "full" mission planned sorties per MR crew	Average number of aircraft not scheduled any activity per week		

The remainder of this section describes in more detail the specific rule and policy changes made to the DOSS baseline model for both aircrew and aircraft scheduling. Later sections compare the effects of these changes on the alternative schedules that result at different levels of sorties and alert.

CHANGES TO THE DOSS BASELINE MODEL

Changes to Aircrew Scheduling

The major changes made to the DOSS baseline model for aircrew scheduling are of two types. The first type involves increasing aircrew availability to fly. This involves (1) policy changes that decrease the time crews spend in mission planning, and (2) rule changes that increase the availability of crews to fly the day before alert and the day after Combat Crew Rest and Recuperation (CCRR) following an alert. The second type of change to the baseline model involves an alternative policy for distributing sorties to MR crews. Changes are made to the DOSS baseline model in order to provide differential training to MR crews based on experience levels.

Flying Before and After Alert. The alert duty is normally seven days of continuous duty where crews are away from their families. Hq SAC policy allows crews 12 hours of crew rest prior to reporting for alert duty (SACR 55-43). Also, Hq SAC policy requires that crews receive at least 50 percent of the total time spent on alert as CCRR (Combat Crew Rest and Recuperation) following an alert. As might be expected, crews feel that flight before alert or immediately after CCRR is undesirable, because it tends to stretch out the period of continuous duty and cuts back on additional crew rest. At both Wing A and Wing B, which were investigated in detail for this study, rules were followed that increased the minimum crew rest period prior to alert, and the CCRR period following an alert. This behavior represents additional resources or slack (in terms of aircrew availability) being expended in order to pay attention to the needs of

aircrews, and thereby to reduce potential conflicts within operations [1]. Uncertainty avoidance is a motive in this area as well [2].

For example, the Hq SAC policy that allows crews 12 hours of crew rest prior to reporting for alert duty, requires crews to be on the ground by 2000 the previous day in order to stand alert by 0800 the next morning (the typical alert changeover time). Therefore, under these restrictions crews are available for flying afternoon CSS sorties the day prior to going on alert, since landing and debriefing is usually completed by 2000. Nevertheless, at both Wing A and Wing B, crews were restricted from flying CSS sorties prior to going on alert. This restriction essentially increases crew rest prior to going on alert.

With regard to flying after CCRR (Combat Crew Rest and Recuperation following an alert duty), Hq SAC policy requires that crews receive at least 50 percent of the total time spent on alert as CCRR. Normally, CCRR is taken for three and one-half days immediately following a seven-day alert duty. This consideration, along with a required 12 hours of crew rest prior to any sortie, allows a crew to fly a morning sortie (takeoff around 0800) immediately following the CCRR period [3]. Mission planning for the sortie would have to be performed while on alert, which is allowable. However, at both bomb Wing A and Wing B, crews were restricted from flying morning sorties following CCRR. This restriction essentially increases the CCRR period following alert.

It should be stressed that the presence of slack resources is not necessarily inefficient. It may be worth the costs to pay attention to the interests of particular subgroups of the organization. The aim

of this study is to make clear what these opportunity costs are, in terms of the effects on a number of output measures when less slack exists. For example, what output is achievable when crews are allowed to fly sorties before alert and immediately after CCRR?

To increase aircrew availability, I change rules in the baseline DOSS model to allow crews to fly afternoon CSS sorties prior to going on alert, and to fly sorties immediately following a CCRR period. These changes do not mean that crews scheduled for alert, or those crews just coming off CCRR, will always be scheduled to these sorties. The changes only allow crews to be scheduled to these sorties, if, given the preference rules for how crews are chosen for particular sorties, they are the most preferred crews to be so assigned. The extent to which such assignments occur is an output measure reflecting morale considerations.

Length of Mission Planning. Mission planning for a typical training sortie normally takes place the day prior to the sortie, and is scheduled for the entire eight-hour work day. With the advent of the B-52 Aircrew Continuation Training Test Program, Hq SAC increased aircrew availability to fly by allowing the use of preplanned mission data, or "canned" mission planning, on a limited basis [4]. The use of "canned" mission planning on a limited basis was designed to considerably decrease time devoted to mission planning, and is incorporated in the DOSS baseline model. Essentially, "canned" mission planning allows mission planning to be scheduled for only four hours prior to takeoff.

The DOSS baseline model also allows sufficient time between sorties for crew rest and mission planning in accordance with Hq SAC

policies. Hq SAC policy requires 12 hours of consecutive crew rest, in addition to time spent in mission planning, prior to all sorties. Taking such considerations into account determines rules for how soon after a particular type of sortie a crew can be scheduled for another type sortie.

At Wing A, because "canned" mission planning is allowed prior to afternoon CSS sorties and night profile sorties, crews could be scheduled both back-to-back CSS sorties and back-to-back night profile sorties. A back-to-back CSS sortie occurs, for example, when a crew flies a CSS sortie on Wednesday and then flies another CSS sortie on Thursday. Although back-to-back sorties do not violate Hq SAC policy, since four hours of "canned" mission planning and twelve hours of crew rest can occur between sorties, they nevertheless are undesirable from a crew's point of view [5]. At both Wing A and Wing B, rules were followed that prohibited back-to-back sorties from being scheduled, and therefore the DOSS baseline model also prohibited such sorties. This restriction is another example of resources (in terms of aircrew availability) being expended to pay attention to aircrew considerations and also to avoid uncertainty [6].

In order to further increase aircrew availability to fly, I change rules in the DOSS baseline model to reflect the policy change of allowing the use of "canned" mission planning for all sorties. Along with the policy change to allow "canned" mission planning prior to all sorties, I also change rules to allow a crew to be scheduled back-to-back sorties of all types (i.e., consecutive day, consecutive CSS, and consecutive night sorties). These back-to-back sorties do not violate Hq SAC policy because "canned" mission planning and

required amounts of crew rest can be scheduled between the sorties, with one exception [7].

The policy change of allowing "canned" mission planning for all sorties does not mean that crews will never receive a full detailed mission planning covering a scheduled eight-hour period, which is considered a training item itself. Full mission planning can still be scheduled to crews, after the flying schedule is produced, to those crews with sufficient time available between sorties. The extent to which "full" mission planning can be scheduled is a measure of training output, and should be compared to the extent of "full" mission planning scheduled under the DOSS baseline model.

Furthermore, allowing crews to be scheduled to back-to-back sorties does not imply back-to-back sorties will always be scheduled. In fact, given the preference rules for choosing aircrews to fly, which will be described shortly, crews are not likely to be the most preferred for a particular type of sortie if just the day previously they had been scheduled that same type of sortie. The number of back-to-back sorties which are scheduled is an output measure reflecting morale considerations.

Differential Training. With the DOSS, I examine the feasibility of an alternative policy for distributing sorties to MR crews, designed to improve overall training effectiveness given a fixed number of sorties. The alternative policy is to provide differential training, or in other words, different numbers of sorties to MR crews based on a measure of their relative proficiency. This is in contrast to the current policy of scheduling each crew an equal number of sorties.

The DOSS baseline model for aircrew scheduling incorporates a variety of preference rules to allocate equal numbers of sorties of different types to all MR crews, in accordance with Hq SAC policy. As mentioned earlier, Hq SAC policy requires that during each quarter every MR crew that is fully available is to receive a minimum of seven event training sorties, four profile training sorties, and one pilot pro sortie for a minimum of 12 total sorties. Because of the differences in training content between non-CSS event sorties and CSS sorties, which are described in detail in App. A, schedulers at both Wing A and Wing B expressed the preference to allocate four non-CSS event sorties and three CSS sorties per crew in order to accomplish the seven event training sortie requirements. Also, schedulers expressed the preference of achieving the four profile sortie requirements per crew by scheduling three night profile sorties and one day profile sortie per crew. Again, this was because of differences in training content, and also because of the number of night and day sorties agreed upon with maintenance. In accordance with the above considerations, the DOSS baseline model considers that each MR crew at the beginning of the quarter needs four morning event sorties, three CSS sorties, three night profile sorties, one day profile sortie, and one pilot pro.

As described earlier, the DOSS baseline model uses a vertical scheduling strategy in scheduling sorties to crews. Therefore, sorties are assigned to crews on a day-by-day basis who meet all constraints (e.g., sufficient time must be available for crew rest and mission planning), and have the highest "need" for a particular type of sortie. "Need" is defined as the number of such sorties remaining

until the minimum requirement for that type sortie has been met. For example, a crew which already has been scheduled to three non-CSS event sorties would have a need of one more in order to meet its requirement. DOSS would prefer this crew, everything else being equal, over another crew which had already been scheduled to four or more such sorties.

An alternative to the policy of equal sortie requirements for all MR crews is the policy of having different minimum sortie requirements based on a crew's relative proficiency--in other words, the policy of differential training. In order to explore the feasibility of such a policy, at different sortie and alert levels, I input to the DOSS baseline model different minimum sortie requirements for MR crews based on a measure of the crew's relative experience. All the many baseline constraints and preference rules are not changed; instead, only the definition of sortie "need" is changed for particular crews.

The measures of experience chosen for this examination are the "S" (select crew), "E" (senior crew), and "R" (ready crew) SAC crew designations for MR crews. These designations are a subjective evaluation of an MR crew's experience relative to other crews at the wing, with an "S" crew being the most experienced, and an "R" crew being the least experienced. Table 5 lists the minimum requirements that were chosen for examination at the current or baseline sortie level. For "S" crews the requirements are decreased, for "E" crews they are not changed, and for "R" crews they are increased from current SAC policy. It should also be noted that the minimum sortie requirement for "S" crews is nine, which was the minimum sortie requirement for all crews prior to the B-52 Aircrew Continuation Training Test Program.

Table 5
MINIMUM SORTIE REQUIREMENTS CHOSEN FOR
EXAMINING DIFFERENTIAL TRAINING--
BASELINE SORTIE AND ALERT LEVEL

	"S" Crews	"E" Crews	"R" Crews
Non-CSS event sorties	3	4	5
CSS sorties	2	3	4
Night profile sorties	2	3	4
Day profile sorties	1	1	1
Pilot pros	1	1	1
Total sorties	9	12	15

The logic behind this method of differentiation is that for any level of sorties flown, greater training effectiveness can be achieved by giving more sorties to the crews with the least experience, and fewer sorties to the more experienced crews [8]. Of course, I use in this examination only one measure related to a crew's relative proficiency. I do not suggest that this is the only measure, or even the best measure, of a crew's proficiency. If a policy of differential training appears feasible and desirable, other measures of a crew's proficiency could be developed and reevaluated over time. The aim of this investigation is to explore the feasibility of providing differential training and to demonstrate a method in which it could be implemented by wing schedulers. Given constraints on aircrew availability, is it feasible to give some crews as many as 15

sorties per quarter on a regular basis? Can schedulers, without the need of computer assistance, provide such differentiation? Both of these questions are addressed in the analysis that I present in later sections.

Since the analysis focuses on the extent to which it is possible, given current resources, to increase sortie and/or alert levels, the analysis also considers rule and policy changes that affect aircraft scheduling.

Changes to Aircraft Scheduling

The major changes that are made to the DOSS baseline aircraft scheduling process involve policy changes to increase aircraft availability to fly. Specifically, the changes are to increase the valid preflight period from 72 hours to 96 hours, and to increase the interval between phase inspections from 100 flying hours to 200 flying hours. In addition, I change rules for scheduling aircraft to turnaround sorties, in the attempt to schedule as many turnaround sorties as possible.

It should again be stressed that DOSS provides no answers to how aircraft reliability may be affected over time due to changes in aircraft scheduling, such as increasing inspection intervals. It provides only a methodology for estimating the possible benefits of such changes--for example, the additional sorties and/or alert that can be scheduled as a result of increased aircraft availability. Other means of analysis besides DOSS, such as flying a test program under alternative conditions, would be needed to completely evaluate the worth of any such changes.

Increasing the Valid Preflight Period. As noted earlier [9], an aircraft is required to have a pre-flight inspection prior to most sorties. Also, after most sorties an aircraft is required to have a basic postflight (BPO) inspection. However, some sorties are classified as turnaround sorties, in which one sortie can follow another with no BPO or preflight in between.

The valid preflight period is that period of time since an aircraft's last preflight inspection, in which the aircraft is allowed to fly a turnaround sortie and therefore not require a new preflight inspection. Increasing this period from 72 to 96 hours is designed to allow for more turnaround sorties, and therefore to decrease preflight and BPO inspections for a given number of sorties. Decreasing the time aircraft spend in preflight and BPO inspections increases aircraft availability to fly, and also reduces maintenance costs in terms of maintenance resources expended in such inspections.

Increasing the Phase Interval. As noted previously, a phase inspection is a scheduled maintenance inspection that typically consumes an entire work week (five working days) that the aircraft is not available to fly. Each phase inspection is actually only one part, or phase, of an overall inspection process. Each phase requires inspecting different systems or components of the aircraft. Currently, one phase inspection occurs at least every 100 flying hours for each B-52 aircraft.

Different scheduled maintenance policies, which affect both the interval and content of phase inspections, have been the subject of much research for many different aircraft of the U.S. Air Force (Cohen, 1972; Elwell, 1976). Determining optimal inspection

policies--involving both interval and content--is a difficult problem. No formal decisionmaking methods exist for reaching optimal inspection content and interval decisions. Scheduled inspection policies tend to be based largely on tradition, experience, and intuition (Cohen, 1972).

The specific policy change I examine with DOSS involves increasing the phase interval for B-52 aircraft from 100 flying hours to 200 flying hours, without any repackaging of the work content of each inspection. Therefore, the length of the phase inspection (five working days) remains the same.

Increasing the phase interval from 100 flying hours to 200 flying hours is designed to decrease the number of phase inspections performed, and thus to increase aircraft availability to fly. Again, I examine such a policy change with DOSS in order to estimate the effects the resulting increased aircraft availability would have on the ability to schedule more sorties and/or alert, given current aircraft resources.

Increasing Turnarounds. As described in Section III, the DOSS baseline model was designed to incorporate the stated preferences of maintenance schedulers for scheduling turnaround sorties. I attempt, through rule changes to the DOSS baseline model, to increase as much as possible the number of turnaround sorties scheduled each week. The number of possible turnaround sorties is limited by the 72-hour pre-flight policy, and by the rule of not flying an aircraft with less than 12 hours ground time since its last flight.

Increasing the number of total turnarounds has several predictable effects. First, scheduling more turnaround sorties each

week reduces the number of aircraft needed to fly a given number of sorties because one aircraft, instead of two, is used to fly two sorties. Using fewer aircraft "frees up" additional aircraft (i.e., increases the number not scheduled to any activity) for other activities, such as increased sorties. Second, more turnarounds decrease the number of preflight and BPO inspections required, as discussed earlier. Therefore, more turnarounds increase aircraft availability to fly, and also reduce maintenance costs in terms of maintenance resources expended in preflight and BPO inspections. Finally, the extent to which more turnarounds can be scheduled is an important measure of maintenance flexibility in substituting aircraft for other scheduled aircraft. For example, if an aircraft cannot fly its scheduled sortie, one alternative to cancelling the sortie is to substitute another aircraft that has recently flown.

Because of the different effects of increasing the number of turnaround sorties, I attempt to schedule with DOSS as many turnarounds as possible, within constraints mentioned above, in order to estimate: (1) the extent to which aircraft can be "freed up" for additional activity; (2) the extent of maintenance costs which potentially would be reduced as a result of fewer preflight and BPO inspections; and (3) the extent of scheduling flexibility that exists to make aircraft or "tail number" substitutions.

I examine the above-mentioned rule and policy changes to the DOSS baseline model for both aircrew and aircraft scheduling, either individually or in combination. Furthermore, the effects of these changes are examined at each of three different sortie and alert levels: (1) the current or baseline level of sorties and alert, (2)

the baseline level of sorties but with increased alert, and (3) increased sortie and alert levels. In summary, the purpose of such an experimental design is to estimate the feasibility and consequences of achieving (a) increased sortie levels, (b) increased alert levels, and (c) significant differential training, both with and without the changes to the DOSS baseline model.

In what immediately follows, I present the analysis of the major effects of changes to the DOSS baseline model, at the baseline level of sorties and alert.

NOTES--SECTION IV

1. Berman (1975:75) also observed similar slack resources being expended, regarding flying before and after alert, at several SAC bomb wings in order to reduce conflict within operations.
2. Schedulers at both Wings A and B expressed that during bad-weather months they avoid flying crews prior to alert because of the risk that they might not return to the home station on time. Berman (1975:75) has also observed this behavior at other SAC bomb wings.
3. At both Wing A and Wing B, alert changeover is on Thursdays at 0800. CRR lasts for three and one-half days until Sunday evening at 2000. The next scheduled sortie is on Monday mornings at 0800.
4. A "canned" mission package allows mission planning to be completed in approximately three hours, reduced from approximately six hours with full mission planning. The reduction is due to the use of pre-planned mission data that reduces a substantial amount of repetitive paperwork involved in the mission planning process. At Wing A "canned" mission planning was performed prior to all CSS and night profile sorties. At Wing B "canned" mission planning prior to sorties was even more constrained, effectively occurring only for CSS sorties.
5. For example, consider the feasibility of back-to-back CSS sorties. Assume takeoff for the first sortie occurs at 1400, landing occurs at 1900, and crew debriefing is

completed by 2000, which is true for the typical scheduled CSS sortie at Wing A. Twelve hours of crew rest beginning at 2000 would expire at 0800 the next day. Therefore, the required crew rest following the first sortie would not interfere with the required four hours of mission planning prior to the second sortie, because the mission planning would normally begin around 1000 on the day of the second CSS sortie.

6. If the first sortie is delayed in landing for some period of time, either because of a late takeoff or because of in-flight problems, the crew may not be able to land in time to receive a full 12 hours of crew rest plus four hours of mission planning prior to the second scheduled sortie. However, a certain amount of slack is also incorporated in the four hours of "canned" mission planning that is typically scheduled. Usually, this type of mission planning is completed in 2.5 to 3.0 hours.
7. DOSS is constrained from scheduling four hours of "canned" mission planning prior to day sorties with a takeoff at around 0800, since this would require crews to report very early in the morning. In adhering to this morale consideration, the four hours of "canned" mission planning is assumed to occur either in the morning or afternoon hours of the day before the flight. For a crew to be scheduled back-to-back day sorties, that crew must be available the entire eight-hour work day prior to the first sortie, in order to mission plan for both sorties in two separate four-hour sessions on the same day. Mission planning for two sorties on the same day is allowed under current SAC policies.
8. This argument assumes that the marginal utility, or training effectiveness, of an additional sortie is greatest for the crew with the least experience.
9. See note 14 of Section III.

V. ANALYSIS AT THE BASELINE SORTIE AND ALERT LEVEL

This section presents an analysis of the effects of rule and policy changes made to the DOSS baseline model. The effects of these changes, on both aircrew and aircraft scheduling, are examined for the case where DOSS is constrained to schedule the baseline level of sorties and alert. As discussed previously, the baseline sortie and alert level refers to what was scheduled on a weekly basis at an actual SAC bomb wing (Wing A).

The rule and policy changes, which were described in Section IV, primarily increase aircrew and aircraft availability. However, because the present analysis is constrained to the baseline level of sorties and alert, this section does not address the increased sorties or alert that may be possible because of increased aircrew and aircraft availability. These issues will be dealt with shortly in Sections VI and VII. Nevertheless, at the baseline sortie and alert level, the rule and policy changes do result in some potential improvements in both aircrew and aircraft scheduling. These improvements are due largely to the increased scheduling flexibility that increased aircrew and aircraft availability allows.

EFFECTS OF CHANGES TO AIRCREW SCHEDULING

The rule and policy changes affecting aircrew scheduling, discussed in Section IV (specifically regarding flying before alert, flying after CCRR, and shortened mission planning), were designed to increase aircrew availability to fly. When these changes are combined, they increase aircrew availability to fly by approximately

45 percent over the baseline case [1]. However, this figure by itself is not very useful for decisionmaking regarding the worth of such changes. What is useful, is to know how this increase in aircrew availability can be used to improve wing effectiveness.

Increases in aircrew availability will increase scheduling flexibility in choosing aircrews for particular sorties. Increased scheduling flexibility, in turn, should result in improved distributions of sorties to mission-ready (MR) crews, regardless of the criterion for distributing sorties. That is, improved distributions should result whether the criterion for distributing sorties is (1) to allocate sorties equally to all crews, or (2) to provide differential training based on experience. Furthermore, because of the different training content of different types of sorties, improved distributions of sorties should also lead to improved distributions of training items.

In addition to potential benefits to training, rule and policy changes that increase aircrew availability may also have some negative aspects that must be considered. For instance, morale may be negatively affected by increasing the number of flights the day before alert.

In this subsection I examine the effects of rule and policy changes to aircrew scheduling on training output, and on measures reflecting morale considerations. I also examine the extent of any differences between what is scheduled and eventually accomplished that might be affected by any of the rule or policy changes.

Effects on Training

Figure 1 below depicts the effects increased aircrew availability has on distributing sorties to MR crews, when the criterion is to allocate sorties equally to all crews in accordance with Hq SAC policy. In order to assess the effects of increased aircrew availability, I compare sortie distributions produced by two DOSS processes: (1) the DOSS baseline process, and (2) the DOSS process that incorporates rule and policy changes to increase aircrew availability. The sortie distributions that both DOSS processes scheduled are also compared to the sortie distributions that were scheduled at Wing A for the same time period. This comparison is made in order to point out an important difference in the way sorties are distributed with DOSS and how they are distributed at the wing.

First, I compare the outputs of the two DOSS processes. The DOSS process that incorporates rules to increase aircrew availability achieved, as expected, an improvement in the distribution of total sorties. This DOSS process scheduled a slightly less varied distribution than the DOSS baseline process scheduled. The standard deviation of total sorties is reduced from 1.09 to .87, due to increased aircrew availability and the consequent increase in scheduling flexibility. Again, the criterion for allocating sorties to crews in this case was to give equal numbers of sorties to all MR crews, which was the expressed objective of wing schedulers in accordance with Hq SAC policy.

Of course, the objective of both DOSS processes, in accordance with Hq SAC policy and the expressed objectives of wing schedulers, was not only to allocate equal numbers of total sorties to all crews,

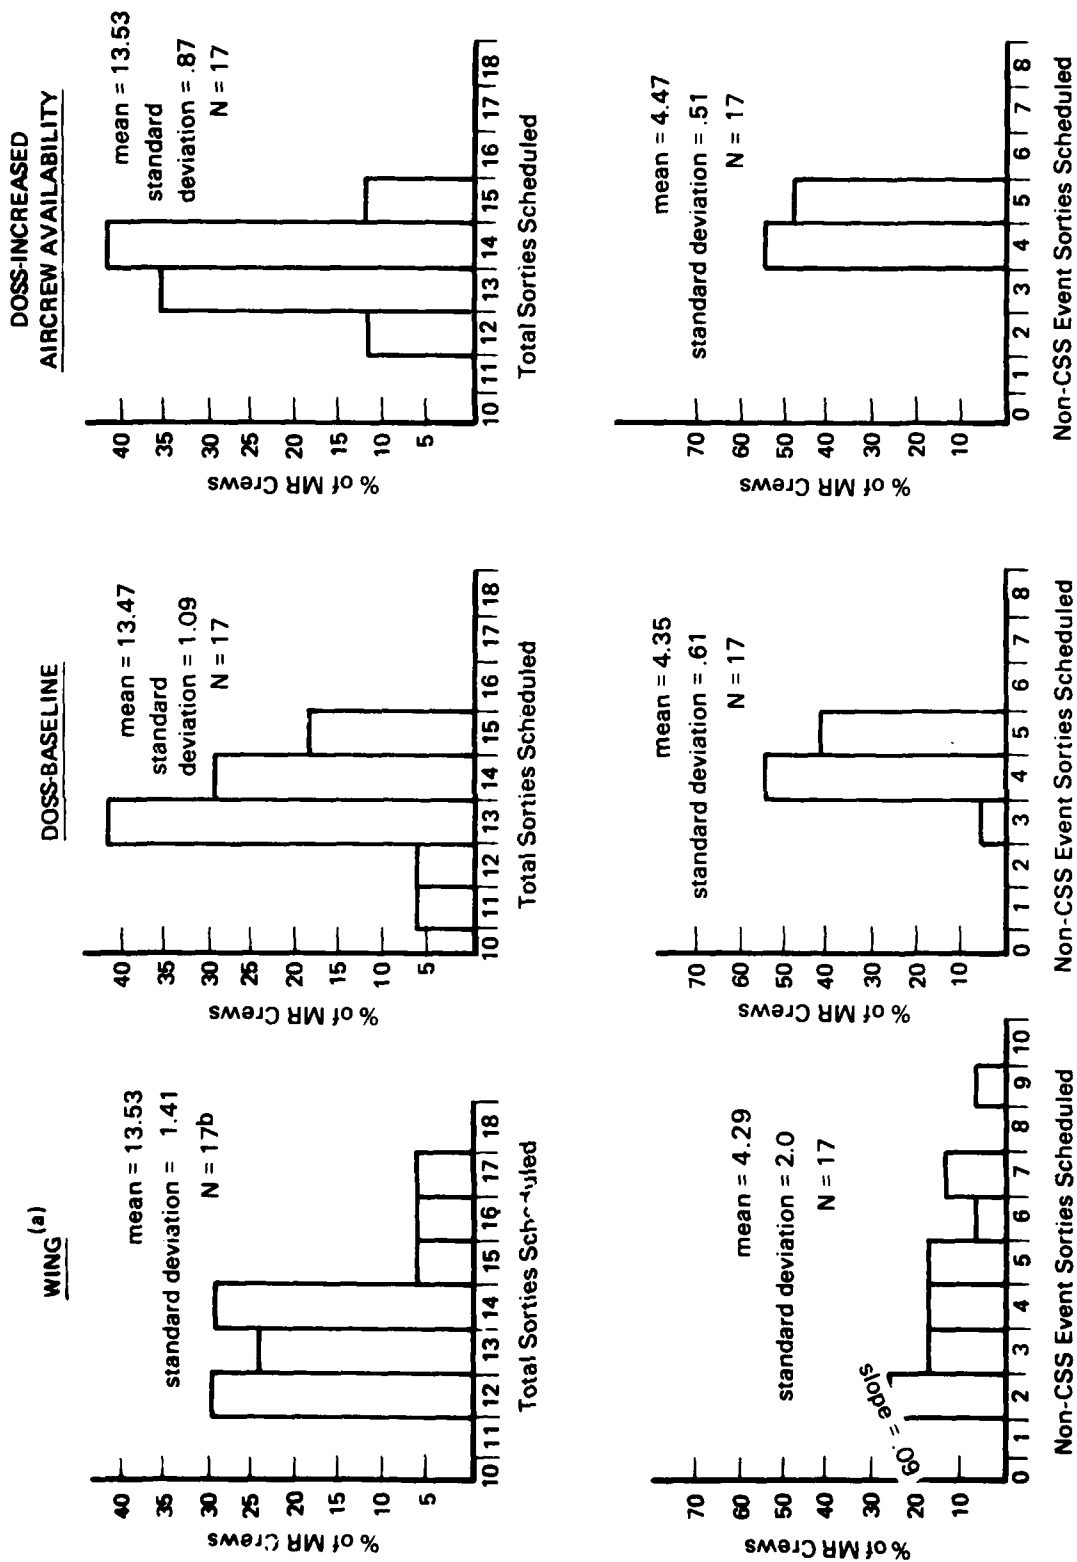


Fig. 1--Mission Ready Crew scheduled sortie distributions (Sept.-Nov. 1978)

- (a) Source: Wing weekly schedules
(b) Considers only crews available for the entire three months

but also to allocate equal numbers of different types of sorties. To illustrate the extent that this objective was achieved with both DOSS processes, Fig. 1 also compares the distributions of one type of sortie; non-CSS event sorties. In this case a slight improvement results from rule and policy changes to increase aircrew availability. The standard deviation of non-CSS event sorties is reduced from .61 to .51 as a result of the increased aircrew availability.

Although sortie distributions that DOSS scheduled were improved as a result of increasing aircrew availability, nevertheless, what is of particular significance is the difference in distributions that the DOSS processes scheduled compared to the distributions scheduled at the wing. In the case of total sortie distributions, wing schedulers produced a more varied distribution than the DOSS baseline process scheduled under identical conditions (standard deviation is 1.41 as compared to 1.09 with DOSS). Even more dramatically, the performance of wing schedulers falls far short of what DOSS showed was possible in attempting to equalize sorties of different types to all MR crews. For example, the standard deviation of non-CSS event sorties scheduled at the wing was 2.0, as compared to a standard deviation of .61 that the DOSS baseline process scheduled under identical conditions. At the wing, one crew was scheduled as many as nine non-CSS event sorties, while four others were scheduled as few as two, even though the expressed objective of wing schedulers was to give a minimum of four such sorties to all MR crews in accordance with Hq SAC policy. Similar wide ranges were scheduled for the other types of sorties as well. Table 6 below compares the ranges of the different types of sorties that were scheduled at the wing, to the ranges that the two DOSS processes scheduled.

Table 6
SCHEDULED SORTIE RANGES (September-November 1978)

	a Wing	DOSS-- Baseline	DOSS--Increased Aircrew Availability
Non-CSS event Sorties	2 - 9	3 - 5	4 - 5
CSS Sorties	1 - 8	2 - 4	2 - 3
Night Profile Sorties	1 - 5	2 - 5	3 - 4
Day Profile Sorties	0 - 4	0 - 4	1 - 3
Pilot Pros	0 - 2	0 - 2	0 - 1
Total Sorties	12 - 17	11 - 15	12 - 15

^a
Based on weekly schedules of Wing A.

Differences in sortie distribution are important because of the differences in training content of the various types of sorties. Widely varied distributions of sorties of different types allocated to MR crews lead to widely varied training item accomplishments. For example, the standard deviation of expected bomb runs (B01s) scheduled per MR crew is reduced from 2.8 (with a range of 20 to 29) under the wing-produced schedules to 1.4 (with a range of 23 to 27) under the DOSS-produced schedule with increased aircrew availability. (For details of the differences in training content of different sorties, and how training item distributions are calculated, see App. A.)

The difference in sortie distributions, between what the wing scheduled compared to what DOSS scheduled, raises an important question about the ability of schedulers to achieve a desired

distribution of sorties (regardless of the criterion for distribution) using current scheduling strategies and without computer assistance. For instance, we have observed with the DOSS baseline process, that the potential exists to provide a much better distribution of training, in terms of equalizing sorties, than wing schedulers provided. The concept of differential training based on experience levels, which I will discuss presently, may further complicate scheduling to some degree, and hence the question arises: Can schedulers, using current scheduling strategies and without the aid of the computer, provide the extent of differential training shown possible with DOSS? If not, can schedulers, without computer assistance, adopt a different scheduling strategy in order to approximate results shown possible with DOSS? This issue is dealt with in detail in Section VIII. The potential for differential training at the baseline sortie and alert level, and at other output levels as well, must first be addressed.

I use two different DOSS processes to examine the potential for differential training at the baseline sortie and alert level. One DOSS process incorporates only those rule changes necessary to schedule differential training, and maintains all other baseline rules. The other DOSS process combines the rule changes necessary to schedule differential training, with the other rule and policy changes previously discussed to increase aircrew availability. The outputs from both DOSS processes indicate that (1) significant differential training is achievable at the baseline level of sorties and alert, and (2) that by increasing aircrew availability the extent of differential training can be improved.

Figure 2 below depicts in detail the extent of differential training based on total sorties that DOSS showed was feasible under the two different conditions previously discussed. Furthermore, the wing-produced distribution of total sorties is also shown as a function of differential sortie requirements, in order to depict the current lack of differential training at the wing. Of course, a lack of differential training is to be expected at the wing since current SAC policy dictates equal sortie requirements for all crews.

The graph in the upper left-hand corner shows what was actually scheduled at the wing. The vertical axis depicts total sorties scheduled, and the horizontal axis depicts different sortie requirements that were given to each crew based on experience level (as shown earlier in Table 5, the "S" crews were given a requirement of 9, the "E" crews 12, and the "R" crews 15). Each dot represents one of 17 crews that were available during the entire three-month period.

The graph of what was scheduled at the wing indicates crews were not allocated significantly different numbers of sorties based on experience levels. A desirable allocation in terms of differential training would be to have sorties scheduled equal to sorties required, with any excess sorties above minimum requirements distributed equally. This was the objective that DOSS sought in attempting to provide differential training. Such a distribution would be represented graphically by "sorties scheduled" as a function of "sortie requirements" being parallel to a line having unit slope (i.e., the dashed line in the graph). We see in this graph, however, that the sorties scheduled are more or less evenly distributed about a

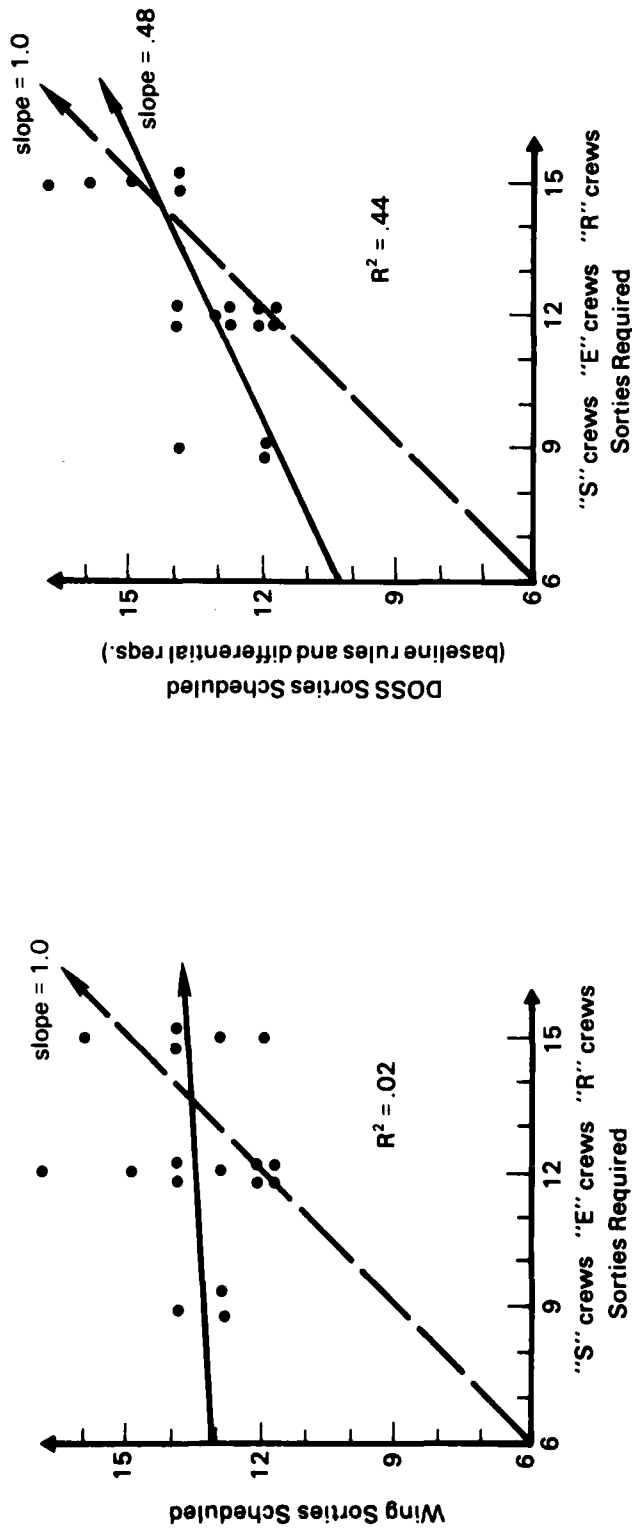


Fig. 2--Potential for differential training at the baseline sortie and alert level*

* For details of the regression equations see Appendix B.

nearly horizontal line (i.e., the solid line in the graph), where sorties scheduled equal approximately 13.5 (the average for all crews). This line, which "best" fits the data points, was determined by an Ordinary Least Squares (OLS) regression (details of which are reported in App. B), where "sorties required" was the independent variable, and "sorties scheduled" the dependent variable. The slope of the estimated OLS line is nearly zero (.09), with sortie requirements explaining only two percent ($R\text{-squared} = .02$) of the variance in sorties scheduled. In other words, different sortie requirements apparently had no significant effect on sortie distribution [2].

The graph in the upper right-hand corner, however, indicates that the DOSS process, which incorporates the baseline rules and is programmed to differentiate on the basis of experience levels, achieved a significant amount of differentiation. In this graph we see that sorties scheduled are more closely correlated to sorties required than was the case at the wing. The OLS line has a greater slope (slope = .48), and sortie requirements explain a higher percentage of the variance in sorties scheduled ($R\text{-squared} = .44$) [3].

Furthermore, the graph at the bottom indicates how much better differential training is made possible by combining the policy changes previously discussed to increase aircrew availability (specifically flying before alert, flying after CCRR, and allowing "canned" mission planning prior to all sorties). In this case, sorties scheduled are highly correlated with sorties required. The slope of the OLS line is .79, and sortie requirements explain as much as 77 percent of the variance in sorties scheduled ($R\text{-squared} = .77$) [4]. This improvement

in differential training is due to the increased flexibility in choosing crews for particular sorties that increased aircrew availability allows.

In addition to differentiating total sortie requirements based on experience levels, the two DOSS processes that attempt to provide differential training also differentiate the sortie requirements of different types of sorties. For instance, "S" crews were given the requirement of three non-CSS event sorties, "E" crews four, and "R" crews five (see Table 5). Therefore, both DOSS processes also attempt to provide differential distributions of sorties of different types based on experience levels. Outputs from both DOSS processes indicate that (1) significant differentiation of sorties of different types is feasible at the baseline level of sorties and alert, and (2) that by increasing aircrew availability the extent of this differentiation can be improved [5].

Because of differences in training content of different types of sorties, differential distributions of sorties lead to differential distributions of training items. For example, in the case of the DOSS-produced schedule that incorporates increased aircrew availability, the estimated number of bomb runs (B01s) scheduled per crew varied significantly given different experience levels. The average number of B01s scheduled per "S" crew was 17.7, for "E" crews was 22.9, and for "R" crews was 30.6. In contrast, at the wing, where differential training was not attempted, the average number of B01s scheduled per "S" crew was 24.0, for "E" crews was 24.0, and for "R" crews was 26.0. (Again, for details of the differences in training content of different sorties, and how training item distributions were calculated, see App. A.)

One training output measure that may be negatively affected by changes to the DOSS baseline model, in order to increase aircrew availability, is the extent of "full" mission planning performed per crew. One of the changes that I examine to increase aircrew availability is the policy change to allow the use of "canned" mission planning for all sorties. Although this policy change increases aircrew availability, it nevertheless may reduce the number of times crews do a "full" mission planning, which is considered a training item itself.

Table 7 below depicts the effects of the policy change to allow "canned" mission planning prior to all sorties, on the extent of "full" mission planning that could still be scheduled to crews. The table shows the average number of "full" mission planned sorties that could be scheduled per crew, as a result of the two DOSS processes that incorporate the policy change. Also depicted is the average number of "full" mission planned sorties each crew was scheduled under baseline rules (i.e., without the policy change).

Essentially, at the baseline sortie and alert level each crew would be scheduled on the average about one less "full" mission planned sortie than would be scheduled without the policy change. The reason for such a slight reduction is due to the fact that although "canned" mission planning is allowed for all sorties, it is not necessary in all cases. Adequate time would still exist between many scheduled sorties in order to schedule "full" mission planning to crews.

Table 7

EXTENT OF "FULL" MISSION PLANNING--
BASELINE SORTIE AND ALERT LEVEL

	DOSS- Baseline	DOSS- Increased Aircrew Availability	DOSS- Increased Aircrew Availability & Differential Training
Average number of "full" mission planned sorties per crew per quarter	7.5	6.5	6.1

Morale Considerations

The many changes made to aircrew scheduling not only affect training, but they also affect several morale considerations. Table 8 below depicts the morale considerations that are affected by the rule and policy changes to increase aircrew availability and to provide differential training.

The two DOSS processes in which rule and policy changes were combined to increase aircrew availability resulted in increasing the average number of back-to-back sorties, flights before alert, and flights after CCRR that were scheduled per MR crew. This was to be expected since the changes to increase aircrew availability were essentially to allow such sorties to be scheduled. Crews, understandably, do not prefer to fly any of these types of sorties. Nevertheless, the number of such sorties scheduled to each crew, given the rule and policy changes, was not large. In fact, as Table 8

shows, the average number of any of these types of sorties that were scheduled per crew did not exceed even one sortie. Furthermore, no one crew was scheduled more than two of any of these types of sorties during the entire period. In addition, these two DOSS processes also result in a reduction in the average workweek of MR crews, from approximately 55 to 52.5 hours [6]. This reduction is due to the decreased time spent in mission planning that results from the policy change to increase the use of "canned" mission planning.

Table 8 also shows that the average duty workweek is less for "S" crews in all cases. This is due to the fact that "S" crews are required to spend less time on alert in order that they have sufficient availability to perform in-flight evaluations of other crews [7]. The DOSS processes that provide differential training result in workweeks that vary even more depending upon a crew's relative experience. This variation results because of different numbers of sorties scheduled to different crews based on their relative experience. "S" crew members are scheduled the fewest work hours under the policy of differential training, and therefore would have additional availability to perform in-flight evaluations of other MR crews. Scheduling such evaluations is currently one of the most difficult problems schedulers face. It should also be noted that the workweek for "R" crew members, in both cases when differential training was provided, is still well below the limit set by Hq SAC (SACR 55-43) of maintaining an average duty workweek over a six-month period below a 74-hour maximum.

Table 8

MORALE CONSIDERATIONS--BASELINE SORTIE AND ALERT LEVEL

Morale Output Measures Per Crew (a)	Wing (b)	DOSS- Baseline	DOSS- Increased Aircraft Availability	DOSS- Differential Training	DOSS- Increased Aircraft Availability & Differential Training
Back-back sorties (c)	0	0	.47	0	.71
Flights be- fore alert (d)	0	0	.06	0	.12
Flights after CCRR (e)	0	0	.65	0	.59
Average Duty Workweek (hours):					
Per MR crew member (f)	55.0	54.9	52.6	54.7	52.4
Per "S" crew member	46.4	46.8	43.9	43.2	41.3
Per "E" crew member	57.0	56.6	54.3	55.6	53.4
Per "R" crew member	56.7	56.7	54.6	60.0	57.1

a
n=17; crews available entire period only (Sept.-Nov. 1978).

b
Based on weekly schedules of Wing A.

c
Refers to consecutive day, afternoon, or night sorties by
the same crew.

d
Refers to flying afternoon CSS sorties the day prior to alert.

e
Refers to flying morning sorties immediately after CCRR.

f
A weighted average of "S", "E", and "R" crews.

Scheduled Versus Accomplished Activity

Up to this point the analysis has focused solely on scheduled activity. Of course, not everything that is scheduled is eventually accomplished. I now address the extent of differences between what is scheduled and eventually accomplished regarding aircrew scheduling. The intent is to assess the effects any such differences would have on the analysis of the alternative rules and policies previously discussed.

There are three major types of sortie cancellations: (1) those due to operations (crews are not available because of illness or emergency leave); (2) those due to maintenance (aircraft are not ready due to system malfunctions); and (3) those due to weather cancellations. In addition, there is a certain amount of training loss on sorties that are successfully flown, due to either weather or aircraft system malfunctions preventing the accomplishment of scheduled training items.

During the period under investigation, September through November 1978, there were ten overall sortie cancellations at Wing A based on weekly schedules. Four cancellations were due to operations, five due to maintenance, and one due to weather [8]. Furthermore, one cancellation due to operations was made up at a later date for a total net cancellation of nine sorties. That is, 239 sorties were originally scheduled for approximately 1377 flying hours, while 230 actually flew for approximately 1328 flying hours. These cancellations resulted in an overall sortie accomplishment rate of .96.

This small difference between what is scheduled and accomplished on a weekly basis has no significant effect on the analysis previously

discussed. Given no maintenance changes to the baseline rules that may decrease the accomplishment rate, I assume that the .96 accomplishment rate applies to the alternatives examined with DOSS. Given the 241 sorties that the DOSS baseline process scheduled, I assume 231 will be accomplished, for a reduction in flying hours from 1385 to 1329. Also, assuming sortie cancellations are evenly distributed across crews, no significant effect on the distributions of sorties to MR crews will result [9]. In addition, given the rule and policy changes that were examined, aircrew availability was always at least as great as with the baseline rules. In fact, it was substantially increased in some cases. This argues that operations schedulers would have at least as much flexibility, under the alternatives that were examined, for rescheduling crews to additional sorties to make up for previous losses--if they so desired [10].

EFFECTS OF CHANGES TO AIRCRAFT SCHEDULING

In addition to examining the effects of alternative rules and policies regarding aircrew scheduling, this study also focuses on alternative rules and policies for aircraft scheduling. The major alternatives chosen for examination were primarily designed to increase aircraft availability to fly. However, because the present analysis is constrained to the baseline level of sorties and alert, this subsection does not address the increased sorties and/or alert that may be possible as a result of increased aircraft availability. These issues are dealt with in Sections VI and VII. Nevertheless, at the baseline sortie and alert level, the changes in aircraft scheduling that were examined do suggest some potential gains. For

AD-A085 394

RAND CORP SANTA MONICA CA

RULE-BASED MODELING AS AN ANALYSIS TOOL: IMPLICATIONS FOR RESOU--ETC(U)

APR 80 R FALLON

F49620-77-C-0023

F/G 5/1

NL

UNCLASSIFIED

RAND/N-1489-AF


2 of 2
AD
UNCLASSIFIED

END
DATE
FILMED
7-80
DTIC

OF

AD

A085394



instance, it may be possible to generate the baseline level of sorties and alert with less maintenance resources (i.e., manpower and aircraft). These resources could then be utilized in other ways besides increasing sorties or alert, such as for additional maintenance on aircraft.

Effects on Maintenance

The major changes affecting aircraft scheduling, described earlier in Section IV, primarily involve policy changes to increase aircraft availability. Specifically, the changes are to increase the valid preflight period from 72 hours to 96 hours, and to increase the interval between phase inspections from 100 flying hours to 200 flying hours. In addition, I change rules for scheduling aircraft to turnaround sorties so as to attempt to schedule as many turnarounds as possible. I examine the effects of these changes independently. That is, I compare the outputs of three different DOSS processes, each of which incorporates only one of the above-mentioned changes to the baseline rules. Table 9 below depicts the effects on several maintenance output measures that result from the different DOSS processes. The table compares these output measures to the output that results from using only baseline rules (i.e., the outputs of the DOSS baseline process for aircraft scheduling). The purpose is to depict the major effects on various maintenance output measures that result because of changes to the baseline rules.

Before discussing the major differences shown in the table, a brief description of the output measures is in order. Preflight inspections, BPOs, phase inspections, and thru-flight inspections

(associated with a turnaround sortie) are maintenance activities that are directly involved with the flying activity. Therefore, the number of these scheduled activities, and the aircraft hours spent in them, are indirect measures of maintenance costs associated with all scheduled flying activity [11]. The average ground time spent by aircraft between scheduled sorties is an indirect measure of the probability of meeting the maintenance schedule. Maintenance crews use these hours to (1) perform various types of inspections in order to prepare the aircraft for its next sortie, and (2) to clear up any "in-flight discrepancies" noted on the previous flight which would prevent the aircraft from flying on its next scheduled sortie [12]. The average number of aircraft scheduled to fly each week, and the average number not scheduled for any activity, are measures of how aircraft are utilized to fly a particular number of sorties [13]. The more turnaround sorties that are scheduled in a week result in fewer aircraft required to fly a given number of sorties. Therefore, more turnarounds essentially "free up" more aircraft (i.e., increase the number of aircraft not scheduled to any activity) for other activities. In addition, the extent to which more aircraft can be "freed up" is an important measure of maintenance scheduling flexibility, to make aircraft substitutions for other aircraft that cannot meet their scheduled sorties.

Table 9

MAINTENANCE OUTPUT MEASURES--BASELINE SORTIE AND
ALERT LEVEL

Maintenance Output Measures	Baseline Rules	96 Hour Valid Preflight	200 Hour Phase Interval	Increased Use of Turnarounds
^a Preflights	106	97	104	81
Turnarounds	86	95	88	111
BPOs	94	85	96	69
Phase inspections	13	13	7	13
Aircraft hours in scheduled inspections	2424	2352	2200	2224
Average hours ground time between sorties	29.1	28.8	29.2	25.2
Average number of aircraft scheduled to fly per week	7.46	7.46	7.46	5.85
Average number of aircraft not scheduled to any activity per week	.62	.62	1.15	2.15

^a
Sept.-Nov. 1978 excluding two sorties scheduled on 9/01; 13
complete weeks.

Being constrained to the baseline sortie and alert level, for the time being, Table 9 does not depict the increased sorties, or aircraft on alert, that may be possible due to increased aircraft availability. These issues will be discussed shortly. However, at the baseline level of sorties and alert some significant differences do result from the changes made to the baseline rules.

As discussed earlier, increasing the valid preflight period is designed to decrease preflight and BPO inspections for a given number of sorties. This, in turn, increases aircraft availability to fly and reduces maintenance costs in terms of resources used in such inspections. For example, in earlier work (Fallon 1979), it was observed that the policy of increasing this period from 48 to 72 hours under the B-52 Aircrew Continuation Training Test Program resulted in a 23 percent reduction in BPOs from what would have been the case under a 48-hour policy. As Table 9 shows, increasing this period further to 96 hours results in a further reduction in BPOs on the order of 10 percent (from 94 to 85). The reduction in preflights and BPOs, of course, must be offset by an increase in turnaround sorties, which are usually accompanied by a thru-flight inspection. Since a thru-flight inspection is typically scheduled for six hours, and a preflight and BPO inspection combine for 14 hours (six scheduled hours for the preflight and eight scheduled hours for the BPO) [14], the reduction in preflights and BPOs should result in less aircraft hours spent in scheduled inspections (preflights, BPOs, thru-flights, and phase inspections). As Table 9 shows, the aircraft hours spent in scheduled inspections are reduced by approximately three percent (from 2424 to 2352 hours). In addition, the increase in turnarounds, which have less ground time on the average between sorties, results in a slight decrease in the average ground time between sorties per aircraft (from 29.1 to 28.8 hours).

Increasing the phase interval from 100 flying hours to 200 flying hours decreases the number of phase inspections required by approximately 50 percent. In the three-month period under

examination, only seven phase inspections were required, as opposed to 13 (one every week) without the policy change. This, naturally, increases aircraft availability to fly. As Table 9 shows, the aircraft hours spent on the ground in scheduled inspections are reduced by over 200 hours (from 2424 to 2200 hours). Also, with this policy option the average number of aircraft not scheduled to any activity per week is increased (from .62 to 1.15). As mentioned earlier, this is an indication of aircraft available for other activities such as more sorties, and is also a measure of scheduling flexibility to make aircraft substitutions.

Lastly, rule changes were made in order to schedule as many turnaround sorties as possible. The number of possible turnaround sorties is limited by the 72-hour preflight policy, and the rule of not flying an aircraft with less than 12 hours ground time since its last sortie. The increased use of turnarounds was attempted in order to estimate: (1) the extent of maintenance costs that potentially could be reduced as a result of fewer preflight and BPO inspections; (2) the extent to which aircraft could be "freed up" for additional activity; and (3) the extent of scheduling flexibility to make aircraft substitutions. As shown in Table 9, the effects of increasing turnarounds were substantial. Preflights and BPOs were reduced by approximately 25 percent, which resulted in an eight percent decrease in hours spent in scheduled inspections (from 2424 to 2224 hours). Also, only five or six aircraft were needed each week in order to meet the schedule, as opposed to seven or eight under baseline rules. This resulted in over two aircraft each week being free of any scheduled activities. Because of the increase in

turnarounds, however, the average ground time between sorties per aircraft was reduced from an average of 29.1 to 25.2 hours.

In all three of these alternatives, hours spent in scheduled inspections were reduced. Also, in the case of increasing turnarounds, the average ground time between sorties per aircraft was substantially reduced. Decreased maintenance, and decreased ground time between sorties, may increase sortie cancellations. For example, decreased maintenance could lead to less reliable aircraft systems over time. Also, decreased ground time between sorties might result in not enough time to fix any problems noted from previous sorties. The extent to which there would be less reliable aircraft systems and increased sortie cancellations can be assessed adequately only by other means of analysis besides DOSS, such as by flying a test program for a sufficient period of time under alternative conditions. The purpose of DOSS is not to answer such questions, which it cannot, but to point out the kinds of benefits that may result under alternative conditions so that more informed judgment can be made regarding all relevant tradeoffs.

Because the present analysis was constrained to the baseline sortie and alert level, the potential for increasing sorties and/or alert was not considered. Nevertheless, at the baseline sortie and alert level the alternatives that were examined did show some potential gains. For instance, we observed that by increasing turnaround sorties it was possible to schedule the baseline level of sorties and alert with fewer aircraft, and with fewer required maintenance activities such as preflight and BPO inspections. This result indicates that a potential exists to achieve current sortie and

alert output with fewer maintenance resources (i.e., manpower and aircraft). These resources could be used in other maintenance activities, such as clearing up delayed discrepancies [15], or could be used solely for increasing scheduling flexibility. Of course, this result also indicates that, with current resources, a potential exists for increasing the level of sorties and alert, which is the focus of later sections.

As in the case of aircrew scheduling, not everything that is scheduled is eventually accomplished. Therefore, the extent of any difference between what is scheduled and eventually accomplished, with regards to aircraft scheduling, must be considered in examining the effects of any rule or policy changes.

Scheduled Versus Accomplished Activity

As noted earlier, at Wing A there were five sortie cancellations due to maintenance during the three-month period under investigation. With regard to aircraft scheduling, however, overall sortie cancellations are not the whole story. Along with the five overall sortie cancellations due to maintenance, there were seven other aircraft malfunctions of various types, each of which prevented a scheduled aircraft from flying its next scheduled sortie. These sorties, along with others in the week the aircraft may have been scheduled for but could not fly because it was not repaired in time, were nevertheless flown by substitute aircraft.

In accordance with Hq SAC guidance (SACR 60-9), maintenance schedulers attempt to substitute, for aircraft that are not able to meet their scheduled sorties, aircraft previously scheduled in the

week. In other words, maintenance schedulers prefer to use turnaround sorties in order to make aircraft substitutions. Therefore, the extent to which it is possible to perform more turnarounds is a measure of flexibility to make aircraft substitutions so that additional sorties will not be cancelled.

Of course, another way to make aircraft substitutions so that additional sorties will not be cancelled is to use aircraft for substitutions that have not been previously scheduled to any activity. Therefore, the average number of aircraft not scheduled to any activity per week is also an important measure of scheduling flexibility. As expected, this measure is highly related to the level of turnarounds scheduled. The more turnarounds that are scheduled, the fewer aircraft that are required to fly a given number of sorties. Therefore, more turnarounds lead to more aircraft not scheduled to any activity.

At the baseline level of sorties and alert, no additional sortie cancellations are to be expected from any of the policy changes considered, since the ability to make aircraft substitutions is not decreased. The 96-hour valid preflight rule, for example, actually increases the ability to perform more turnaround sorties. This, as previously discussed, increases scheduling flexibility in making aircraft substitutions. Furthermore, the 200-hour phase-interval policy results in more aircraft on the average not scheduled to any activity per week, which also provides more flexibility to make aircraft substitutions. Lastly, the rule changes that result in an increased use of turnarounds, by definition, imply that no further flexibility exists for increasing turnarounds in order to make

aircraft substitutions. However, under this option, Table 9 shows that over two aircraft, on the average, would be "free" each week for making aircraft substitutions. This number of available aircraft is far more than would be needed to make the same number of substitutions as occurred at Wing A during the period under investigation [16].

At additional levels of sorties or alert, where resources are more constrained, sufficient flexibility may or may not exist to make the baseline level of aircraft substitutions. If sufficient flexibility does not exist, this may result in additional sortie cancellations not accounted for in the observed .96 sortie accomplishment rate. The extent of this flexibility, and the extent of any additional sortie cancellations, will be dealt with when additional sortie and alert levels are considered.

SUMMARY

This section focused on the analysis of various rule and policy changes made to the DOSS baseline model for both aircrew and aircraft scheduling. The changes were primarily designed to increase aircrew and aircraft availability to fly. The effects of these changes, on both aircrew and aircraft scheduling, were examined for the case where DOSS was constrained to schedule the baseline level of sorties and alert. Because the analysis was so constrained, it did not address the increased sorties and/or alert that may be possible as a result of increased aircrew and aircraft availability. These issues will be dealt with shortly in Sections VI and VII. Nevertheless, at the baseline level of sorties and alert, the rule and policy changes did result in some potential improvements in both aircrew and aircraft

scheduling. The improvements were due largely to the increased scheduling flexibility that increased aircrew and aircraft availability allowed.

The major finding with regard to aircrew scheduling was that sufficient aircrew availability existed to provide significant differential training based on experience levels. However, with rule and policy changes to increase aircrew availability, the extent of differential training was improved considerably. It is argued that differential training based on a measure of the relative proficiency of crews will provide greater training effectiveness, given any fixed number of sorties.

The major finding with regard to aircraft scheduling was that the baseline level of sorties and alert could be scheduled with fewer aircraft and with fewer required maintenance activities such as pre-flight and BPO inspections. This result indicates that a potential exists to achieve current sortie and alert output with less maintenance resources (i.e., manpower and aircraft). The result also indicates that a potential exists, within current resources, to achieve increased levels of sorties and alert. The extent of this potential is addressed in the following two sections.

NOTES--SECTION V

1. Aircrew availability to fly is calculated on the basis of the number of aircrews available to fly at various takeoff times during the week. For example, only five of 18 MR crews are typically available to fly Monday morning sorties (approximate takeoff time 0800) under the baseline set of rules. However, 10 crews would be available to fly if we account for the rule and policy changes that increase aircrew availability. Similar calculations are made for

each of the other takeoff times during the week. The percentage increases in aircrew availability that result at each takeoff time are then averaged. This approach to calculate aircrew availability is also followed at the other sortie and alert levels that are discussed throughout the text.

2. The t-statistic for the independent variable "sortie requirements" is only .5. With 15 degrees of freedom, the hypothesis that "sortie requirements" have no effect on sortie distribution cannot be rejected at the .05 significance level. For details of this and other regressions reported in this study, see App. B.
3. The t-statistic for the independent variable "sortie requirements" in this case was 3.41. With 15 degrees of freedom, the hypothesis that "sortie requirements" have no effect on sortie distributions can be rejected at the .05 significance level. For details, see Table B1 of App. B.
4. The t-statistic for "sortie requirements" in this case was 6.7. With 15 degrees of freedom, the hypothesis that "sortie requirements" have no effect on sortie distributions can be rejected at the .05 significance level. For details, see Table B1 of App. B.
5. For example, OLS regressions were computed with "non-CSS event sorties scheduled" as the dependent variable and "non-CSS event sortie requirements" as the independent variable. As in the case of total sortie distributions, a desirable differential allocation of non-CSS event sorties in terms of differential training is one where the slope of the estimated OLS line is nearly equal to 1.0, and where a high R-squared statistic results. In the case of the DOSS process programmed to provide differential training under baseline rules, the slope of the OLS line was .52; the R-squared was .34; and the t-statistic for the independent variable "non-CSS event sortie requirement" was 2.8. With 15 degrees of freedom this t-statistic value is significant at the .05 level. In the case of the DOSS process that incorporates increased aircrew availability, the slope was .79; the R-squared was .61; and the t-statistic was 4.8, which again is significant at the .05 level. For details of these regressions, see Table B2 of App. B.
6. I estimated the average duty workweek of aircrews (for the 13-week scheduling period) based on the following three components: (1) time crews spent on alert; (2) time crews spent in activities directly associated with flying; and (3) time spent by the entire crew in ground training activities. In estimating time spent on alert, each day that a crew was on alert counted as 24 hours of continuous duty. In order to estimate the time crews spent in flying activities, the

approximate flying time of each sortie was included, along with time for (a) mission planning (4 or 8 hours), (b) show-up prior to flight (1 hour), and (c) debriefing and report completion following each flight (1 hour).

7. The most senior "S" crew is not to receive more than 60 percent of the average non-"S" crew alert rate, while the remaining two "S" crews are not to receive more than 70 percent (SACR 55-43).
8. These cancellations do not include numerous CSS cancellations and later additions due to exceeding the one-and-one-half-hour maximum ground time for CSS sorties. They also do not include many sortie cancellations and later additions due to flying the back-up aircraft for the CSS sorties and not the primary aircraft.
9. Assuming that a significant relationship exists between differential sortie requirements and sorties scheduled, and that a .96 sortie accomplishment rate applies to all crews, implies that a significant relationship will also exist between sortie requirements and sorties accomplished. All coefficients in the regressions are simply multiplied by a constant term of .96. Also, the t-statistic for the independent variable "sortie requirements" is not affected. Therefore, the question of whether or not to reject the hypothesis that a significant relationship exists at the .05 level is not affected.
10. As noted earlier, only once was a crew rescheduled a sortie to make up for a previous loss. When operations schedulers were asked about this they indicated that usually no crew rescheduling takes place, since a certain amount of attrition is planned for and accepted.
11. I estimated aircraft hours spent in scheduled inspections based upon the following parameters:
 - a. 6 hours per preflight inspection
 - b. 8 hours per BPO inspection
 - c. 6 hours per thru-flight inspection
 - d. 40 hours per phase inspection (5 work days).

These parameters represent the time that was typically scheduled for each of the inspections at both Wing A and Wing B.

12. The average hours of ground time spent by aircraft between sorties is calculated on the basis of the following averages:

- a. 18 hours per quick-turnaround sortie (average ground time of all turnarounds occurring in less than 24 hours).
- b. 32 hours per normal-turnaround sortie (average ground time of all turnarounds occurring in more than 24 hours).
- c. 40 hours for sorties requiring preflight inspections.

For example, the DOSS baseline process scheduled 106 pre-flights, 83 normal-turnarounds, and 3 quick-turnarounds, for a total of approximately 6,950 hours that all the aircraft spent on the ground between sorties $[(106 \times 40) + (83 \times 32) + (18 \times 3)]$. Given 239 sorties were scheduled, the average ground time between sorties is 29.1 hours. The average of 1.5 hours ground time between CSS sorties was not included, since virtually no maintenance is performed during this time.

- 13. For example, in the baseline case there were 8 of 13 weeks in which one aircraft was not scheduled any activity--an average of .62 aircraft per week.
- 14. See note 11.
- 15. "Delayed discrepancies" refer to various aircraft-related problems that require maintenance attention, but for which the needed work has been postponed or delayed.
- 16. The seven situations where aircraft substitutions were required at Wing A each occurred in separate weeks. One previously scheduled aircraft in each of these separate weeks was not able to fly all or some of its sorties scheduled for that week. Therefore, having at least one aircraft not scheduled to any activity the entire week, for each week of the quarter, would ensure enough availability to make the same level of aircraft substitutions as occurred at Wing A.

VI. ANALYSIS AT INCREASED ALERT

One option for increasing the capability of the U.S. strategic bomber force is to increase the percentage of aircrews and aircraft on continuous ground alert. Currently, approximately 30 percent of the bomber force is on continuous ground alert (U.S. Congress, 1979:397). The policy option of increasing the strategic bomber alert force involves many concerns that are well beyond the scope of this study. In evaluating such a policy option, the strategic nuclear capability of the U.S. relative to the Soviet Union, along with the future role of our strategic bomber force, are issues of primary concern. However, increasing the percentage of aircrews and aircraft on alert has several training, maintenance, and morale implications, which cannot be overlooked in any thorough assessment of such a policy option.

When an aircrew or aircraft is on alert it is not available to fly training sorties. Therefore, a tradeoff exists between the size of the alert force and the number of training sorties that can be scheduled. This section presents an analysis of the effects of rule and policy changes to the DOSS baseline model, for the case where DOSS attempts to schedule the baseline level of sorties, yet with the alert force increased from four to five aircrews and aircraft on alert. The purpose is to examine the feasibility and the effects (on a number of training, maintenance, and morale measures) of maintaining increased alert, both with and without the changes made to the DOSS baseline model.

EFFECTS OF CHANGES TO AIRCREW SCHEDULING

Having an additional aircrew on alert each week constrains aircrew availability to fly. In fact, such an increase in the alert level, without changing the level of sorties attempted, decreases aircrew availability to fly by approximately 30 percent [1]. However, even with this increased alert, aircrew availability can actually be increased by 18 percent over the baseline case, with the combined rule and policy changes discussed earlier (specifically, changes regarding flying before alert, flying after CCRR, and allowing the use of "canned" mission planning prior to all sorties). In the present analysis I use DOSS in order to examine the effects of these changes in aircrew availability on a number of training and morale measures.

Specifically, I use DOSS to attempt to schedule the baseline level of sorties but with increased alert, under two different sets of conditions. The first involves changes to the DOSS baseline process for aircrew scheduling to allow for differential training, while maintaining all other baseline rules. The second set of conditions involve rule and policy changes to allow for increased aircrew availability, combined with changes to provide differential training. I examine the outputs of these two DOSS processes on various measures of training and morale. I also examine the extent of any differences between what is scheduled and likely to be accomplished under the alternative conditions.

Effects on Training

A major finding of the analysis was that the increased alert level severely constrained aircrew availability to fly. The DOSS process that maintained all baseline rules, yet attempted to provide differential training, resulted in a six percent reduction in total sorties, coupled with a much poorer differential training distribution. Fourteen sorties could not be scheduled because of insufficient aircrew availability. Furthermore, no significant relationship existed between total sortie requirements based on experience levels and total sorties scheduled. The slope of the estimated OLS line in this case was only .20, with sortie requirements explaining only eight percent of the variance in total sorties scheduled ($R\text{-squared} = .08$) [2]. Not surprisingly, then, there was also no significant relationship between the requirements for different types of sorties, and the sorties of different types scheduled [3]. These results were due entirely to the 30 percent reduction in aircrew availability that resulted from increased alert, since at the baseline sortie and alert level DOSS scheduled all sorties as well as significant differential training.

However, the DOSS process that incorporated the combined rule and policy changes to increase aircrew availability successfully scheduled the baseline sortie level along with the increased alert. Furthermore, a significant amount of differential training was provided as well. In this case, the slope of the estimated OLS line was .86, with sortie requirements explaining 86 percent of the variance in total sorties scheduled ($R\text{-squared} = .86$) [4]. In addition, DOSS achieved significant differentiation of sorties of

different training content [5], which as observed earlier leads to significant differentiation of training items. These improvements in training output at this increased alert level were due entirely to the changes adopted to achieve the 18 percent increase in aircrew availability. This increase in aircrew availability resulted in increased scheduling flexibility for choosing particular crews for particular sorties.

The extent of "full" mission planning that could be scheduled was negatively affected by increasing the alert level, and the policy change to allow the use of "canned" mission planning for all sorties. This policy change, along with the increased alert level, resulted in a reduction in the average number of "full" mission planned sorties that could be scheduled per crew, from 7.5 under the DOSS baseline case (as shown above in Table 7) to 5.5.

Morale Considerations

Table 10 below depicts output measures reflecting morale considerations, that result from the DOSS process that schedules the baseline level of sorties and increased alert (allowing for differential training and increased aircrew availability). These measures are compared to the output that results from (1) the DOSS baseline process, and (2) the DOSS process allowing for differential training and increased aircrew availability, yet constrained to schedule the baseline sortie and alert level. The purpose is to depict the major effects on morale considerations that result because of the increased alert and the changes made to the baseline rules.

Table 10

MORALE CONSIDERATIONS

Morale Output Measures Per Crew (a)	Baseline Level of Sorties and Alert		Baseline Level of Sorties and Increased Alert	
	Baseline Rules	Differential Training and Increased Aircrew Availability	Differential Training and Increased Aircrew Availability	
Average crew alert rate (b)	.24	.24	.30	
Back-back alerts	0	0	0	
Back-back sorties(c)	0	.71	1.65	
Flights before alert (d)	0	.12	.29	
Flights after CCRR (e)	0	.59	.65	
Average Duty Workweek (hours):				
Per MR crew member	54.9	52.4	61.7	
Per "S" crew member	46.8	41.3	45.0	
Per "E" crew member	56.6	53.4	64.1	
Per "R" crew member	56.7	57.1	67.3	

a

n=17; crews available entire period only (Sept.-Nov. 1978).

b

Alert rate, as computed here, is the number of weeks on alert divided by the number of weeks in the period; 13. "S" crews are excluded; n = 14.

c

Refers to consecutive day, afternoon, or night sorties by the same crew.

d

Refers to flying afternoon CSS sorties day prior to alert.

e

Refers to flying mornings immediately after CCRR.

In increasing the alert level from four to five aircrews on alert, it was necessary to produce a new alert schedule, taking as given all other activities such as leaves and TDYs. As Table 10 shows, this resulted in an increase in the average alert rate, excluding "S" crews, from .24 to .30. The average number of alerts per crew was increased from 3.14 per quarter to 3.93. A very significant negative morale factor is the back-to-back alert--where a crew goes on alert only one week since its last alert. Crews dislike back-to-back alerts, and schedulers attempt to avoid them as much as possible. With five crews on alert there was no need to schedule any back-to-back alerts. At any additional alert levels, however, back-to-back alerts would be necessary given the same total number of crews [6].

The number of back-to-back sorties was increased as a result of increasing the alert level. At the baseline alert level the number of total back-to-back sorties was 12, or .71 per crew, when incorporating rule changes to allow for such sorties. Under identical rules at the increased alert level, the total back-to-back sorties were 28, or 1.65 per crew. Such an increase is a negative morale factor associated with increasing alert. However, the increase resulted in no more than three back-to-back sorties scheduled to any one crew for the three-month scheduling period. The most significant negative morale factor that resulted was the increase in workweek for all three categories of MR crews, primarily due to the increased alert rate. However, even at this alert level, workweek for "R" crews is still below the limit established by Hq SAC of maintaining an average duty workweek over a six-month period below a 74-hour maximum.

Scheduled Versus Accomplished Activity

The small difference between what is scheduled and accomplished with regard to aircrew scheduling has no significant effect on the analysis of aircrew scheduling at the increased alert level. As was the case at the baseline alert level, if we assume that the .96 accomplishment rate of sorties is evenly distributed across all crews, then no significant effect on the distribution of sorties to crews results. Also, with the rule and policy changes to increase aircrew availability (which we saw were needed in order to meet the baseline sortie level), there is an 18 percent increase in aircrew availability over the baseline case. Thus, even more flexibility exists to reschedule crews than existed under the baseline case.

The analysis pertaining to aircrew scheduling has demonstrated the feasibility of increasing alert, while still maintaining the baseline level of sorties (provided specific rule and policy changes to increase aircrew availability are adopted). It remains to be seen, however, whether sufficient aircraft availability exists to achieve both the baseline sortie level and the increased alert.

EFFECTS OF CHANGES TO AIRCRAFT SCHEDULING

Effects on Maintenance

Table 11 below depicts the major effects, on various maintenance output measures, of increasing the alert level from four aircraft on alert to five. The table compares the output of several different DOSS processes for aircraft scheduling that attempt to schedule the baseline level of sorties but with increased alert. The differences

in the DOSS processes are due to the separate rule and policy changes made to the baseline rules (specifically, the 96-hour valid preflight policy, the 200-hour phase interval policy, and increased use of turnarounds). These changes are primarily designed to increase aircraft availability to fly. The outputs of these DOSS processes are also compared to the output from the baseline case (i.e., the DOSS baseline process for aircraft scheduling, which adheres to baseline rules at the baseline sortie and alert level). The purpose is to depict the major effects on various maintenance output measures that result because of the increased alert and the changes to the baseline rules.

Table 11

MAINTENANCE OUTPUT MEASURES

Maintenance Output Measures	Baseline Level of Sorties and Alert				
	Baseline Rules	Baseline Rules	96-Hour Valid Preflight	200-Hour Phase Interval	Increased Use of Turnarounds
Total Sorties ^a	239	239	239	239	239
Preflights	106	99	93	100	81
Turnarounds	86	93	99	92	111
BPOs	94	87	81	93	69
Phase Inspections	13	13	13	7	13
Aircraft Hours in Scheduled Inspections	2424	2368	2320	2176	2224
Average Hours Ground Time Between Sorties	29.1	28.2	28.0	28.5	25.2
Average Number of Aircraft Scheduled to Fly per Week	7.46	7.08	7.08	7.08	5.85
Average Number of Aircraft Not Scheduled Any Activity per Week	.62	.23	.23	.77	1.15

^a
Sept.-Nov. excluding two sorties scheduled on 9/01; 13 complete weeks.

The most significant result shown in Table 11 is the fact that the baseline level of sorties, along with the increased alert, could be scheduled without any rule or policy changes to increase aircraft availability. That is, the DOSS baseline process was able to schedule the baseline level of sorties, even with one additional aircraft on alert (and thus not available to fly). In order to do so, however, the number of turnarounds was increased eight percent, from 86 to 93. This increase in turnarounds, along with a seven percent decrease in preflight and BPO inspections, resulted in a two percent decrease in total aircraft hours spent in scheduled inspections (from 2424 to 2368 hours). The increase in turnarounds also resulted in a slight decrease in average ground time between sorties from 29.1 to 28.2 hours.

The 96-hour valid preflight policy further decreased preflights and BPOs, and therefore increased aircraft availability to fly, as it was designed to do. However, this increased availability was not necessary to schedule the baseline level of sorties and increased alert.

The 200-hour phase interval policy reduced the required number of phase inspections, and thereby increased aircraft availability to fly. But again, this increased aircraft availability was not necessary to schedule the baseline level of sorties and increased alert. Nevertheless, this policy change did increase the average number of aircraft not scheduled to any activity per week from .23 to .77. These resources can be used for other activities, such as for additional sorties, or simply as increased scheduling flexibility to make aircraft substitutions.

Scheduled Versus Accomplished Activity

As discussed in detail earlier, the capability for making aircraft substitutions, in order to avoid additional sortie cancellations, depends largely upon the extent to which it is possible to perform more turnaround sorties, and in a highly related way on the number of aircraft not scheduled to any activity per week. As Table 11 shows, under the baseline set of rules for aircraft scheduling, the number of aircraft not scheduled to any activity per week decreases from .62 to .23 at the increased alert level. This indicates possibly more difficulty in making the same number of aircraft substitutions as occurred at Wing A. Any increased difficulty in making aircraft substitutions could lead to more sortie cancellations than reflected in the observed .96 accomplishment rate at Wing A. However, even with increased alert, there still is a large potential to perform more turnarounds. In fact, as Table 11 shows, the potential is large enough to "free up" on the average 1.15 aircraft each week, for the entire week, when different rules are followed in order to schedule as many turnarounds as possible. This potential for increased turnarounds is more than is needed to make the same number of substitutions as occurred at Wing A [7]. Therefore, no additional sortie cancellations are to be expected, at the increased alert level, as a result of any difficulty in making aircraft substitutions.

As noted above, without any changes to the baseline rules, an eight percent increase in turnarounds was necessary in order to schedule the baseline sortie level and the increased alert. Whether this increase in turnarounds would have any adverse effect on the observed .96 sortie accomplishment rate, because of possibly lower

aircraft reliability and less time to repair aircraft between sorties, cannot be adequately assessed by DOSS. Other means of analysis are needed to address these types of issues.

SUMMARY

This section focused on the analysis of various rule and policy changes made to the DOSS baseline model for both aircrew and aircraft scheduling. The changes were primarily designed to increase aircrew and aircraft availability to fly. The effects of these changes, on both aircrew and aircraft scheduling, were examined for the case where DOSS attempted to schedule the baseline level of sorties, yet with the alert force increased from four to five aircrews and aircraft on alert. The purpose was to examine the feasibility and the effects (on a number of training, maintenance, and morale measures) of maintaining increased alert, both with and without the changes made to the DOSS baseline model.

The major finding of the analysis with regard to aircrew scheduling was that the baseline level of sorties along with increased alert could be scheduled, given current aircrew resources, provided that certain rule and policy changes to increase aircrew availability are adopted (for example, allowing the use of "canned" mission planning prior to all sorties). Furthermore, the analysis demonstrated that a significant potential for differential training also exists at this level of sorties and alert. The most significant negative factor that resulted from increasing the alert level was an increase in the average workweek for aircrews from approximately 55 to 62 hours.

The major finding of the analysis with regard to aircraft scheduling was that aircraft availability was not a limiting factor in scheduling the baseline level of sorties along with increased alert. DOSS scheduled this level of output without any rule or policy changes to increase aircraft availability. This increase in output, however, necessitated an eight percent increase in turnarounds, and a similar decrease in preflight and BPO inspections.

NOTES--SECTION VI

1. For a description of how aircrew availability is calculated, see footnote [1] of Section V.
2. The t-statistic for the independent variable "sortie requirements" in this case was only 1.15, which is not significant at the .05 level with 15 degrees of freedom. See Table B1 of App. B for the details of this regression.
3. For example, when "non-CSS event sorties scheduled" were regressed on "non-CSS event sortie requirements," the R-squared was only .14, with a t-statistic for "non-CSS event sortie requirements" of only 1.57. With 15 degrees of freedom this t-statistic is not significant at the .05 level. See Table B2 of App. B for details of this regression.
4. The t-statistic for "sortie requirements" in this case was 9.6, which is significant at the .05 level with 15 degrees of freedom. See Table B1 of App. B for details of this regression.
5. For example, when "non-CSS event sorties scheduled" were regressed on "non-CSS event sortie requirements," the R-squared was .32 with a t-statistic for "non-CSS event sortie requirements" of 2.7, which is significant at the .05 level. See Table B2 of App. B for details of this regression.
6. For example, when the case of six aircrews on alert was examined, it was necessary to schedule 17 back-to-back alerts; one per MR crew. For an overview of the analysis performed with DOSS at this increased alert level, see Fallon (1979).
7. See note 16, Section V.

VII. ANALYSIS AT INCREASED SORTIE AND ALERT LEVELS

The analysis undertaken with DOSS in this study was primarily designed to estimate the extent of any additional capability within SAC bomb wings to increase output in terms of sorties and alert. For this reason, the effects (on a number of training, maintenance, and morale measures) of increasing the alert level was examined in Section VI. The purpose of this section is to estimate whether additional capability exists to increase output levels even further. To this end, the present analysis focuses on the effects of rule and policy changes to the DOSS baseline model, for the case where DOSS attempts to schedule both increased alert and sortie levels.

In addition to increasing the alert level from four aircrews and aircraft on alert to five, additional changes are made in order to increase sorties as much as possible, given current flying hour constraints. Essentially, I change the mix of total sorties scheduled, by increasing the number of shorter five-hour sorties that are scheduled (non-CSS and CSS event sorties), and decreasing the number of longer seven-hour sorties that are scheduled (day and night profile sorties). However, I allow a minimum number of long sorties to remain in the schedule in order to fulfill HHDM responsibilities, as well as for participation in the diversity program [1]. In addition, I increase the number of CSS sorties (which, as described earlier, involve very little maintenance), in order to keep maintenance costs from increasing significantly. When all of these changes are considered simultaneously, they lead to a possible total sortie increase of 12.5 percent, or 30 more sorties than were

scheduled at the baseline level. I refer to this level of sorties as the maximum level of sorties given current flying hour allocations.

At this increased sortie and alert level the many rule and policy changes previously described to increase aircrew and aircraft availability and to provide differential training are made to the DOSS baseline model, and alternative schedules are produced. The purpose is to examine the feasibility and the effects (on a number of training, maintenance, and morale measures) of achieving this increased output level, given the rule and policy changes made to the DOSS baseline model.

EFFECTS OF CHANGES TO AIRCREW SCHEDULING

Adding an additional aircrew to alert each week, along with attempting to schedule the maximum level of sorties holding total flying hours constant, severely constrains aircrew availability to fly. In fact, aircrew availability to fly is reduced from the baseline case by approximately 49 percent [2]. However, even at this much greater output level, aircrew availability can still be increased by three percent over the baseline case, with the combined rule and policy changes discussed earlier (specifically, changes regarding flying before alert, flying after CCRR, and allowing the use of "canned" mission planning prior to all sorties).

In the present analysis, I use DOSS to attempt to schedule the maximum level of sorties along with increased alert, given rule and policy changes to allow for increased aircrew availability and to provide for differential training. I examine the outputs of this DOSS process for aircrew scheduling on various measures of training and

morale. I also examine the extent of any differences between what is scheduled and likely to be accomplished at this increased sortie and alert level.

Effects on Training

The major finding of the analysis was that the DOSS process incorporating the combined rule and policy changes to increase aircrew availability successfully scheduled the maximum level of sorties, holding flying hours constant along with the increased alert. In this case all 271 sorties that DOSS attempted to schedule were scheduled. Also, a highly significant amount of differential training was achieved at this increased sortie and alert level. The slope of the estimated OLS line was .89, with sortie requirements explaining 82 percent of the variance in sorties scheduled ($R\text{-squared} = .82$) [3]. In addition, DOSS achieved significant differentiation of sorties of different training content [4], which as observed earlier leads to significant differentiation of individual training items.

Increasing the sortie level, holding flying hours constant, affects not only total sorties, but in addition: (1) the mix of sorties of different training content, (2) total training items scheduled, and (3) the diversity of Radar Bomb Scoring (RBS) sites.

Table 12 below compares the mix of sorties scheduled at the maximum sortie level, holding flying hours constant, with the mix that was scheduled at the baseline sortie level. The major differences involve the substantial increase in five-hour event sorties (non-CSS and CSS event sorties) from 126 to 228, and the sharp decline in the longer profile sorties (day and night) from 102 to 29. The 29 longer

sorties, as mentioned earlier, were required in order to fulfill HHDMs and to participate in the diversity program.

Table 12

SORTIES SCHEDULED (SEPTEMBER-NOVEMBER 1978)

	Baseline Sortie Level	Maximum Sortie Level Holding Flying Hours Constant
Total Sorties	241	271
Flying Hours (a)	1383	1385
Non-CSS Event Sorties	79	137
CSS Sorties	47	91
Night Profile Sorties	58	3
Day Profile Sorties	44	26
Separate Pilot Pros (b)	13	14

a

Flying hours are approximated using the averages of five hours per non-CSS event sortie, five hours per CSS sortie, seven hours per profile sortie, and three hours for a pilot pro.

b

Some pilot pros are dual logged on day profile sorties. At the baseline sortie level 6 were scheduled this way for a total of 19 pilot pros; at the increased sortie level 5 were scheduled this way, again for a total of 19 pilot pros.

Total sorties, holding flying hours constant, could be increased even further over what I refer to as the maximum sortie level if only pilot pros are flown. However, given the very limited training value of such sorties, total training items accomplished would decrease

significantly. Rather, when I speak of the maximum sortie level I imply that not only are flying hours held constant but also that total training items are not reduced. In fact, as shown in Table 13 below, total training items scheduled are actually increased (for an important set of items) at the maximum sortie level. This increase is due to the different mix of sorties of different training content that is scheduled at the maximum sortie level. (For details of the differences in training content of different sorties, and how total training items scheduled are estimated, see App. A.)

Table 13

TOTAL TRAINING ITEMS SCHEDULED (a)

Selected Training Items	Baseline Sortie Level	Maximum Sortie Level Holding Flying Hours Constant
Total Sorties	241	271 (+12.5%)
Flying Hours	1383	1385
B01s--Low Altitude Bomb Runs	550	597 (+9%)
N09s--TA/EVS Navigation Legs	219	242 (+11%)
N15s--Low Altitude Navigation Legs	219	242 (+11%)
R01s--Air Refueling	187	^b 187
E01s--RBS ECM Run	503	506 (+1%)
F01s--Fighter Intercepts	26	^c 26

^a
For an explanation of all terms, and how total training items resulting from the different schedules are estimated, see App. A.

^b
More air refuelings could be scheduled, but the baseline level is assumed so that no new tanker resources are required.

^c
More fighter intercepts could be scheduled, but the baseline level is assumed so that no new fighter resources are required.

A negative factor, however, in increasing sorties while holding flying hours constant, is a possible reduction in the diversity of RBS sites. Currently, Hq SAC directs that each MR crew should practice simulated bomb releases on at least six different RBS sites per quarter, in order to be exposed to sufficient bombing diversity.

However, scheduling fewer long-duration sorties may make it difficult to schedule crews to practice simulated bombing at RBS sites that, because of geographic location, require long sortie lengths. At the baseline sortie level there is a sufficient number of long sorties (night and day profile sorties) to enable each crew to visit at least six different RBS sites. However, at the maximum level of sorties, with 18 crews, the number of long sorties is only sufficient to allow seven crews to visit four different sites, and eleven crews to visit five different sites [5]. This possible decrease in RBS site diversity must be weighed against the increases in sorties and training items shown possible at this increased sortie level.

Another factor of importance in terms of training is the reduction in the number of night sorties that results from attempting to schedule the maximum sortie level. This reduction is considered a negative factor since certain training value is placed on night sorties. At the baseline sortie level, all 58 night sorties are profile sorties. At the maximum sortie level, 21 CSS sorties are scheduled at night (details of how and when CSS sorties are scheduled will be discussed shortly), which, when combined with three scheduled night profile sorties, results in a total of 24 night sorties.

Another training output measure is the extent of "full" mission planning that can be scheduled during the three-month scheduling period. At the maximum sortie and increased alert level, 6.9 was the average number of sorties per crew where "full" mission planning could have been scheduled. This is higher than the 5.5 average that resulted under identical conditions at the baseline level of sorties, which we observed earlier in Section VI. This difference is

principally due to the fact that the average number of sorties scheduled per crew is increased at the higher sortie level from 13.53 to 15.12. This increase naturally provides more opportunities for "full" mission planning.

Morale Considerations

Table 14 below depicts output measures reflecting morale considerations that result from the DOSS process that attempts to schedule the maximum level of sorties and increased alert. Relevant comparisons are made to the outputs of DOSS processes incorporating identical rules and policies, yet constrained to schedule at the different levels of sorties and alert previously discussed. The purpose is to depict any major effects on morale considerations that result because of the increased output.

In scheduling the increased sorties and alert, the number of back-to-back sorties, flights before alert, and flights after CCRR per crew are increased, as would be expected, over all previous cases. However, even at this level of sorties and alert only two crews get as many as four back-to-back sorties, no crew gets more than one flight before alert, and no crew receives more than two flights after CCRR. Average duty workweek is also increased for each category of MR crew, yet not significantly over the case where the baseline level of sorties is scheduled with increased alert. This indicates that the major increase in workweek comes from the additional alert duty and not the increase in sorties. Again, however, workweek for "R" crews is still below the limit established by Hq SAC of maintaining an average duty workweek over a six-month period below a 74-hour maximum.

Furthermore, the average scheduled workweek for "S" crews is not significantly different than under the baseline case (45.9 as opposed to 46.8), which implies still sufficient availability to perform in-flight evaluations of other crews even at the increased sortie and alert level.

Table 14

MORALE CONSIDERATIONS

Morale Output Measures Per Crew (a)	Baseline Level of Sorties and Alert		Baseline Level of Sorties and Increased Alert	Maximum Level of Sorties and Increased Alert
	Baseline Rules	Differential Training and Increased Aircrew Availability	Differential Training and Increased Aircrew Availability	Differential Training and Increased Aircrew Availability
Average crew alert rate (b)	.24	.24	.30	.30
Back-back alerts	0	0	0	0
Back-back sorties(c)	0	.71	1.65	2.4
Flights before alert (d)	0	.12	.29	.47
Flights after CCRR (e)	0	.59	.65	1.2
Average Duty Workweek (hours):				
Per MR crew member	54.9	52.4	61.7	62.6
Per "S" crew member	46.8	41.3	45.0	45.9
Per "E" crew member	56.6	53.4	64.1	65.0
Per "R" crew member	56.7	57.1	67.3	68.1

a

n=17; crews available entire period only (Sept.-Nov. 1978).

b

Alert rate, as computed here, is the number of weeks on alert divided by the number of weeks in the period; 13. "S" crews are excluded; n = 14.

c

Refers to consecutive day, afternoon, or night sorties by the same crew.

d

Refers to flying afternoon CSS sorties the day prior to alert.

e

Refers to flying mornings immediately following CCRR.

Scheduled Versus Accomplished Activity

In order to schedule the maximum level of sorties, holding flying hours constant without significantly increasing maintenance costs (in terms of preflights, BPOs, etc.), I increased the number of scheduled CSS sorties from 47 under the baseline sortie level to 91 under the increased sortie level. This was accomplished in a variety of ways, such as flying two CSS sorties given a choice of three input aircraft, and scheduling one night CSS sortie following two afternoon sorties. Different ways of scheduling CSS sorties imply different accomplishment rates. Furthermore, these accomplishment rates can be estimated, given data on the accomplishment of CSS sorties obtained from both wings, which participated in the B-52 Aircrew Continuation Training Test Program [6].

The result of taking these differences into account is that the expected overall accomplishment rate of CSS sorties is reduced from .89 to .87, implying only 79 of 91 CSS sorties scheduled are expected to be flown. This reduced accomplishment rate of CSS sorties leads to an overall sortie accomplishment rate of .94, as compared to the observed .96 accomplishment rate at the baseline sortie level. A .94 sortie accomplishment rate implies that although 271 sorties are scheduled at the maximum sortie level, only 255 would likely be accomplished. This results in a 10.0 percent increase in expected sorties accomplished over the baseline level (from 231 to 255).

This small difference between what is scheduled and likely to be accomplished at the maximum level of sorties and increased alert has no significant effect on the analysis previously discussed regarding aircrew scheduling at this increased output level. Assuming the .94

accomplishment rate is evenly distributed across all crews, no significant effect on the distribution of sorties to crews will result. Also, with the rule and policy changes to increase aircrew availability three percent over the baseline case, approximately the same flexibility would exist to reschedule crews.

EFFECTS OF CHANGES TO AIRCRAFT SCHEDULING

Effects on Maintenance

Table 15 below depicts the major effects on various maintenance output measures that result from attempting to schedule the maximum level of sorties (holding flying hours constant) along with increased alert. The table compares the output of several different DOSS processes for aircraft scheduling that attempt to schedule the maximum level of sorties along with increased alert. The differences in the DOSS processes are due to the separate rule and policy changes made to the baseline rules (specifically, the 96-hour valid preflight policy, the 200-hour phase interval policy, and rule changes to increase the use of turnarounds). These changes were primarily designed to increase aircraft availability to fly. The outputs of the different DOSS processes are also compared to the output from the baseline case (i.e., the DOSS baseline process that adheres to baseline rules at the baseline level of sorties and alert). The purpose is to depict the major effects on various maintenance output measures that result because of the increased sorties and alert, and the changes to the baseline rules.

Table 15

MAINTENANCE OUTPUT MEASURES

Maintenance Output Measures	Baseline Level of Sorties and Alert	Maximum Level of Sorties and Increased Alert			
	Baseline Rules	Baseline Rules	96-Hour Valid Preflight	200-Hour Phase Interval	Increased Use of Turnarounds
a					
Total Sorties	239	270	270	270	270
CSS Sorties	47	91	91	91	91
PreFlights	106	100	94	101	92
Turnarounds	86	79	85	78	87
Quick-turnarounds	3	12	12	10	65
BPOs	94	87	81	94	80
Phase Inspections	13	13	13	7	13
Aircraft Hours in Scheduled Inspections	2424	2284	2236	2106	2234
Average Hours Ground Time Between Sorties	29.1	23.5	23.3	23.7	20.6
Average Number of Aircraft Scheduled to Fly per Week	7.46	7.0	7.0	7.0	6.85
Average Number of Aircraft Not Scheduled Any Activity per Week	.62	.15	.15	.7	.31

^a
Sept.-Nov. 1978 excluding scheduled sorties on 9/01; 13 complete weeks.

The most significant result shown in Table 15 is the fact that the increased level of sorties that DOSS attempted to schedule, along with increased alert, could be scheduled without any changes to increase aircraft availability. Aircraft availability does not appear constraining even at this greater level of sorties and alert. Furthermore, without any changes to the baseline rules, preflights and BPOs were reduced by approximately seven percent from what occurred at the baseline level of sorties and alert. Furthermore, turnarounds were also reduced by eight percent (from 86 to 79). These reductions were offset by the increase in CSS sorties from 47 to 91, which resulted in a six percent reduction in aircraft hours spent in scheduled inspections (from 2424 to 2284 hours). The increase in total sorties and CSS sorties is also reflected by the decrease in average ground time between sorties from 29.1 to 23.5 hours.

The 96-hour valid preflight policy further decreased preflights and BPOs, and therefore increased aircraft availability, as it was designed to do. However, this increase in availability was not necessary to schedule the increase in sorties and alert.

The 200-hour phase interval policy reduced the required number of phase inspections, and thereby increased aircraft availability to fly. But again, this increase in availability was not necessary to achieve the increased output. Nevertheless, this policy change did increase the average number of aircraft not scheduled to any activity per week from .15 to .7, which can be used to increase scheduling flexibility to make aircraft substitutions.

Scheduled Versus Accomplished Activity

As discussed earlier, the capability for making aircraft substitutions in order to avoid sortie cancellations depends largely upon the extent to which it is possible to perform more turnaround sorties, and in a highly related way on the number of aircraft not scheduled to any activity per week. As Table 15 shows, at the maximum level of sorties and with increased alert, the number of aircraft not scheduled to any activity per week decreases from .62 to .15 under the baseline sets of rules. This suggests that possibly more cancellations than are reflected in the expected .94 accomplishment rate may result due to difficulties in making aircraft substitutions. Also, the potential for performing more turnarounds does not appear large enough to "free up" a substantially greater number of aircraft to be used for substitutions. For instance, when attempting to schedule as many turnarounds as possible at the maximum sortie level, the number of turnarounds goes up, as expected, but only by 10 percent from the number of turnarounds DOSS scheduled when following only baseline rules (from 79 to 87). The average number of aircraft not scheduled to any activity per week was thereby increased from .15 to only .31. This implies that, even when turnarounds are increased as much as possible, only in 4 out of the 13 weeks will an aircraft be available the entire week to be used as a substitute aircraft. This is not enough potential aircraft availability to ensure that the same number of substitutions that occurred at Wing A could be made [7]. However, having an aircraft available the entire week for substitution purposes is a very conservative measure of whether all substitutions could be made. Having sufficient potential to perform more

turnarounds so that one or more aircraft could be "freed up" for only part of a week may be all that is needed.

In order to estimate whether sufficient potential exists at the increased level of sorties and alert to make the same number of aircraft substitutions as occurred at Wing A, I attempted to reschedule in order to make the necessary substitutions. I rescheduled for seven separate weeks, where aircraft failures requiring aircraft substitutions were assumed to take place [8]. The result was that the same number of substitutions that occurred at Wing A could be made, thereby resulting in no additional sortie cancellations not reflected in the .94 accomplishment rate.

As noted above, without any changes to the baseline rules, decreased maintenance is scheduled at the increased level of sorties and alert due to more CSS sorties and fewer preflights and BPOs. This decrease was reflected by the six percent reduction in aircraft hours spent in inspections. Also, the average ground time per sortie was decreased significantly (from 29.1 to 23.5 hours). Again, whether these changes would have any adverse effect on the expected sortie accomplishment rate, because of possible lower aircraft reliability and less time to repair aircraft between sorties, cannot be adequately assessed by DOSS. Other means of analysis are needed to address these types of issues.

SUMMARY

This section focused on the analysis of various rule and policy changes made to the DOSS baseline model for both aircrew and aircraft scheduling. The changes were primarily designed to increase aircrew and aircraft availability to fly. The effects of these changes, on both aircrew and aircraft scheduling, were examined for the case where DOSS attempted to schedule both increased sorties and alert. The purpose was to examine the feasibility and the effects (on a number of training, maintenance, and morale measures) of achieving this level of increased output, given the changes made to the DOSS baseline model.

The major finding of the analysis with regard to aircrew scheduling, was that significant increases in sorties and alert could be scheduled, given current aircrew resources, provided that certain rule and policy changes to increase aircrew availability are adopted (for example, allowing the use of "canned" mission planning prior to all sorties). Furthermore, the analysis demonstrated that a significant potential for differential training also exists at this increased level of sorties and alert. The most significant negative factors that resulted from increasing the level of sorties and alert were: (1) a possible reduction in RBS site diversity, and (2) an increase in the average workweek for aircrews.

The major finding of the analysis with regard to aircraft scheduling was that aircraft availability was not a limiting factor in scheduling the increased sorties and alert. DOSS scheduled the increased output without any rule or policy changes to increase aircraft availability. Furthermore, because of increased use of CSS sorties, the increase in output was accomplished with fewer preflight,

BPO, and thru-flight inspections (associated with turnarounds). At the increased level of sorties and alert, however, only a small potential still existed to achieve this level of output with less aircraft by flying more turnaround sorties. This indicates that any further increases in sorties and/or alert would probably be constrained by aircraft availability.

NOTES--SECTION VII

1. The diversity program refers to sorties where a crew leaves the home base, flies to another bomb wing where crew rest takes place, and then later in the week returns to the home base, thus accomplishing two long-duration profile sorties in the process. These sorties are high in training content, and provide crews the opportunity to practice low-level procedures over diverse terrains.
2. For a description of how aircrew availability is calculated, see footnote [1] of Section V.
3. The t-statistic for the independent variable "sortie requirements" was 8.3, which is significant at the .05 level with 15 degrees of freedom. For details of this regression, see Table B1 of App. B.
4. For example, when "non-CSS event sorties scheduled" were regressed on non-CSS event sortie requirements, the R-squared was .64. The t-statistic for the independent variable "non-CSS event sortie requirements" was 5.2, which is significant at the .05 level with 15 degrees of freedom. For details of this regression, see Table B2 of App. B.
5. The number of RBS sites on which crews from a given bomb wing can practice simulated bombing is a function of many factors, such as the number of RBS sites, the number of sorties scheduled, and the distances of RBS sites from the home base. In order to estimate the effect on bombing diversity that results from alternative sortie levels and mix of sortie types, the following two assumptions were made:
 - o At least three different RBS sites are within a five-hour mission length
 - o At least three different RBS sites are within a seven-hour mission length.

Both bomb wings that were examined in detail in this investigation met the above criteria.

6. The procedure followed at both test wings (Wing A and Wing B) in scheduling CSS sorties was to schedule two aircraft to land at approximately the same time, and then have one of these two aircraft fly the CSS sortie within 1.5 hours after landing. It is important to note, however, that the procedure called for one of the input aircraft to be designated the "primary" aircraft and the other as the "back-up" aircraft. At the time schedulers produce the weekly schedule the "primary" aircraft is the aircraft scheduled to fly the CSS sortie. Since the "primary" aircraft is scheduled to be flown, and subsequent scheduling decisions have been based on this assumption, the primary is the aircraft that will be chosen to fly the CSS sortie, provided that it develops no serious system malfunctions on its previous flight. In other words, the probability of the "primary" aircraft being chosen to fly the CSS sortie over the "back-up" aircraft is independent of the condition of the "back-up" aircraft.

The observed accomplishment rate of CSS sorties at both test wings was approximately .89. Given that the observed probability of flying one CSS from a choice of two input aircraft is .89, and because the probability of choosing the "primary" is essentially independent of the condition of the "back-up" aircraft, the following probabilities of flying CSS sorties given different numbers of input aircraft can be deduced:

- a. .96 probability of flying one CSS sortie given three input aircraft.
 - b. .85 probability of flying two CSS sorties given three input aircraft.
 - c. .67 probability of flying one CSS sortie from only one input aircraft.
7. See note 16 of Section V.
 8. The seven weeks where substitutions took place at Wing A were the same seven weeks where I attempted to reschedule, for the same number and types of sorties substituted for at the wing.

VIII. IMPLEMENTATION ISSUES FOR SCHEDULING

One of the major findings of the DOSS analysis was that a significant potential for differential training exists at all the sortie and alert levels that were examined. However, the existence of the potential for differential training does not guarantee that schedulers can achieve the results shown possible with DOSS. For instance, we observed earlier that the DOSS baseline process for aircrew scheduling provided a much better distribution of training sorties than did wing schedulers, based upon the objective (expressed by operations schedulers and in accordance with current SAC policy) to allocate equal numbers of sorties of different types to all mission ready crews (see Fig. 1 and Table 6). This raises a key issue about the ability of schedulers to achieve a desired distribution of sorties (regardless of the criterion for distribution), given their current scheduling strategies and without computer assistance. The concept of differential training may further complicate scheduling by requiring schedulers to allocate different numbers of different types of sorties to aircrews, based on some measure of a crew's relative proficiency. Hence the question arises: Can schedulers, using current scheduling strategies and without the aid of the computer, provide the extent of differential training that DOSS has shown is feasible?

In order to address this key issue, I managed a scheduling "exercise" in which three SAC operations schedulers from Wing A participated. Each scheduler was asked to produce flying schedules for the identical period and under identical sets of rules and conditions that DOSS was faced with. The purpose was to estimate the

extent of differential training that wing schedulers could provide, under conditions where DOSS has shown significant differential training was possible. Furthermore, the schedulers were asked to follow two distinct approaches or strategies to scheduling for each set of conditions. The purpose was to estimate the effects of these alternative scheduling strategies on the extent of differential training that could be scheduled.

The two different scheduling strategies examined in this exercise were the horizontal and vertical strategies described earlier. The horizontal strategy is the common scheduling strategy that operations schedulers currently follow when producing a quarterly flying schedule. Basically, the scheduler attempts to fill in a matrix with scheduled flying activity where the columns of the matrix are the days of the period and the rows represent distinct crews. The scheduler first schedules flying activity by considering one crew at a time. He moves horizontally across a particular crew's row until he schedules all required training sorties for that crew. In so doing, the scheduler takes into account the many rules, both constraints and preferences, discussed earlier for scheduling sorties to crews--for example, rules that specify whether sufficient time has elapsed since the crew's last sortie to allow for crew rest and mission planning. Given sufficient days available for a crew to fly, it is relatively easy to schedule a crew for close to exactly the total number of sorties desired, and also the desired number of different types of sorties.

The difficulty with this scheduling approach arises near, or at the end, of having scheduled all the crews, or rows of the matrix. At

this time the number of sorties of different types scheduled for a particular day (or vertical column) almost never match with the number and types of sorties contracted for in advance with maintenance that are planned to be generated on that particular day. Because of numerous mismatches between what is scheduled and required, the scheduler then proceeds in a very unsystematic way to make marginal changes to the schedule, wherever possible, in order to match the number of sorties scheduled each day to the requirement. In the process of doing so, however, the number of total sorties, and especially the number of different types of sorties previously scheduled to a crew, will change. This requires further iterations, usually taking several tedious hours, until the scheduler produces what he considers to be a "satisfactory" schedule.

As mentioned earlier, the expressed objective of wing schedulers in accordance with Hq SAC policy is to give equal numbers of sorties to all MR crews, above some minimum of 12. Furthermore, schedulers desire, in accordance with Hq SAC policies, to equally distribute sorties of different training content so as to equally distribute training items that crews accomplish. However, because of the difficulties of equally distributing sorties and still matching daily sortie requirements, schedulers have expressed, and we have observed, a criterion for a "satisfactory" schedule that is much less stringent. As observed through the DOSS analysis, total sorties could have been distributed more equally than were distributed at the wing. Furthermore, the distribution of sorties of different types was widely varied at the wing, which (as the DOSS analysis has shown) did not need to be the case (see Fig. 1 and Table 6).

Providing differential training on the basis of a crew's experience level would not appear to significantly complicate scheduling using a horizontal strategy. The only difference is that crews, or rows, would need to be separated into distinct classes with different "horizontal" requirements (i.e., sortie requirements). Therefore, my hypothesis going into the scheduling exercise was that schedulers, using a horizontal strategy, would find it relatively no more difficult to differentially allocate total sorties as they do in equalizing total sorties to crews. However, difficulties would still persist in scheduling a specific number of sorties of different types.

The vertical scheduling strategy, which the schedulers were asked to apply in the scheduling exercise, is designed along lines similar to how DOSS schedules sorties. This strategy essentially treats each day, or vertical column, independently. The particular daily requirement for different types of sorties, contracted for with maintenance, are scheduled to those available crews which have the most "need" for the sorties. A crew is considered available if it meets a host of constraints, such as whether sufficient time has elapsed since the crew's last sortie to allow for crew rest and mission planning. A crew is considered to have the greatest "need" for a particular type of sortie on the basis of a number of preference rules. For example, the available crew with the highest "need" for a type A sortie will be that crew with the highest number of type A sorties remaining until it fulfills its type A sortie requirement. Over time, sorties of different types should be distributed in a way highly correlated to the specified requirements at the beginning of the training period. If a crew, because of high amounts of alert or

leave, falls behind in terms of sorties flown relative to other crews, then that crew will become the most preferable to fly until it catches up.

It is no more difficult, using the vertical strategy, to differentially allocate sorties of different types, than it is to allocate equal numbers of sorties of different types to all MR crews. The only difference is that the definition of sortie "need" is changed at the beginning of the training period. Also, a routine set of rules is followed systematically for each day of the period. Therefore, no one day is any more difficult to schedule than any other. Furthermore, at the end of the period all daily sortie requirements should be met. In addition, if the routine rules for choosing crews for particular sorties are properly followed, the distribution of sorties of different types should be highly correlated with sortie requirements. We have observed this to be the case with the DOSS-produced schedules.

My hypothesis going into the scheduling exercise was that schedulers, when attempting to follow a vertical strategy, would be able to satisfactorily apply the routine "vertical" scheduling procedures outlined above, without any computer assistance. Furthermore, they would be able to achieve a significant degree of correlation between sortie requirements and sorties scheduled, not only for total sorties but also for sorties of different types. The sortie requirements that schedulers were asked to adhere to were the same as those adhered to by DOSS (i.e., the differential requirements based on experience levels depicted in Table 5).

Scheduler A, the most experienced of the group, attempted to provide differential training at the baseline level of sorties and alert, with all baseline rules held constant. He was asked to do this twice; once using the horizontal strategy he was very familiar with, and secondly using the vertical strategy which had just been explained to him only the hour before. Scheduler B, the least experienced of the group, attempted to provide differential training at the baseline level of sorties but with the number of aircrews on alert increased from four to five. In addition, the rule and policy changes described earlier to increase aircrew availability were allowed. Scheduler B was asked to follow only the vertical strategy. Scheduler C was asked to schedule under the identical conditions and rules facing scheduler B, but to use only the horizontal strategy he was very familiar with.

The results of this exercise were entirely consistent with my hypotheses. Both schedulers A and B were able to satisfactorily follow the vertical strategy, using nothing more than a pad of paper and pencil to keep track of sortie requirements remaining for each crew. Scheduler A, with one and one-half years experience, took approximately 2.5 hours to produce each schedule using both strategies. Scheduler B took 4.5 hours to complete his schedule using the vertical strategy, but this extra time can be attributed to the fact he had never before produced a quarterly schedule. He essentially was a brand new scheduler who was being trained as a replacement for scheduler A. Scheduler C, who had approximately one year of experience, completed his schedule using a horizontal strategy in approximately 3.5 hours.

In order to measure the extent of differentiation of total sorties scheduled in each case, I performed an Ordinary Least Squares regression, where the dependent variable was total sorties scheduled and the independent variable was total sorties required. Similarly, "non-CSS event sorties scheduled" are regressed on "non-CSS event sortie requirements" to illustrate the extent of differentiation scheduled for different types of sorties. The results of these regressions are reported in Table 16 below. As discussed previously, a slope of the estimated OLS line equal to 1.0 represents the case where sorties scheduled are perfectly aligned with sorties required. The R-squared statistic represents the percentage of the variance in sorties scheduled explained by sortie requirements. Furthermore, the t-statistic is an indication of whether the independent variable, sortie requirements, has a statistically significant effect on sorties scheduled. Given that each regression involves 15 degrees of freedom, a t-statistic value greater than 2.13 is statistically significant at the .05 significance level.

Table 16

RESULTS OF SCHEDULING EXERCISE

	Regressions on Total Sorties			Regressions on Non-CSS Event Sorties		
	Slope	R-Squared	t-statistic	Slope	R-Squared	t-statistic
Scheduler A Horizontal Strategy	.73	.64	5.2	.15	.01	.3
Scheduler A Vertical Strategy	.91	.82	8.4	1.2	.54	4.2
Scheduler B Vertical Strategy	.87	.95	16.7	.81	.65	5.2
Scheduler C Horizontal Strategy	.61	.65	5.2	.41	.13	1.5

The results show that scheduler A, when following a horizontal strategy, achieved a significant degree of differentiation of total sorties based on experience levels. However, as expected, he was not able to achieve significant differentiation of sorties of different types. Nevertheless, when following the vertical strategy he achieved highly significant differentiation of total sorties, as well as sorties of different types.

Scheduler C, using the horizontal strategy (but at an increased alert level), did provide significant differentiation of total sorties. Nevertheless, as was the case with scheduler A, scheduler C could not provide significant differentiation with regard to sorties of different types. However, scheduler B provided significant differentiation for sorties of all types. Scheduler B had no prior

experience and was scheduling under conditions identical to those of scheduler C, except for the fact that he was following the vertical strategy.

In sum, this exercise demonstrates that it is feasible for schedulers, without computer assistance, to provide significant differential training at alternative levels of sorties and alert. This is true, however, provided that the commonly employed horizontal scheduling strategy is changed to the vertical strategy described above. It appears that only by following this alternative strategy can a specific criterion for distributing sorties of different training content be adequately applied. This is true whether the criterion is to equally distribute sorties of different training content as is current policy, or to provide differential training based on experience levels. In the case of current policy, the definition of sortie "need" for different types of sorties would simply be the current mission-ready crew sortie requirements. The schedulers who participated in this exercise, seeing the benefits the vertical strategy could provide in terms of better equalizing sorties of all types, have subsequently adopted such a strategy in their everyday scheduling.

IX. CONCLUSIONS AND IMPLICATIONS

This study has served several purposes. The primary aim has been to examine the potential of rule-based modeling as an analysis tool. To this end I used a rule-based model (DOSS) to analyze alternative policies and decision rules that affect resource allocation within a large public-sector organization: the Strategic Air Command. The analysis focused on the policy issues of (a) increasing the SAC alert force, and (b) flying more training sorties, in order to estimate the extent of any additional capability that may exist within SAC bomb wings given current resources. Naturally, the analysis has implications for the particular issues that were examined, but its more general aim was to illustrate the potential of DOSS for examining a broad range of resource allocation policy issues facing SAC, and possibly other Air Force Commands as well.

Because of the different purposes of this study, the intent of this section is threefold:

1. to briefly summarize the major findings and policy implications of the analysis;
2. to discuss the general implications of DOSS for improving resource allocation within SAC; and finally,
3. to discuss some general implications of rule-based modeling for improving organizational decisionmaking.

MAJOR FINDINGS AND POLICY IMPLICATIONS OF THE ANALYSIS

In Sections V through VII, the analysis focused on the effects of alternative decision rules and policies regarding aircrew and aircraft scheduling within a typical SAC bomb wing. The purpose was to estimate the capability of SAC bomb wings to fly more sorties and to increase the alert force. My intent in this subsection is to first summarize the major findings and implications of the analysis pertaining to aircrew scheduling. Second, I will summarize the major findings and implications of the analysis regarding aircraft scheduling.

Major Results of Changes to Aircrew Scheduling

Rule and policy changes were made to the DOSS baseline model for aircrew scheduling in order to (a) significantly increase aircrew availability to fly (such as allowing the use of "canned" mission planning prior to all sorties), and to (b) provide differential training based on experience levels. The effects of these changes were analyzed at each of three different sortie and alert levels: (1) the baseline level of sorties and alert, (2) the baseline level of sorties but with increased alert, and (3) increased sortie and alert levels.

A major finding of the analysis was that increased levels of sorties and alert are possible, provided that certain rule and policy changes to increase aircrew availability are adopted. Furthermore, the analysis demonstrated that a significant potential for differential training exists. As defined earlier, differential training refers to distributing sorties based on a measure of the

relative proficiency of crews, as opposed to giving every crew the same number of sorties. The major implication of differential training is that greater overall training effectiveness can be achieved given any fixed number of sorties.

Besides the positive implications for training and alert that result from (a) more sorties, (b) increased alert, and (c) significant differential training, the analysis also pointed out some negative aspects of increasing the level of sorties and alert. The most serious of these aspects are (a) a possible reduction in RBS site diversity for training purposes, and (b) an increase in the average workweek for aircrews.

Even though the analysis of rule and policy changes to aircrew scheduling demonstrated that additional sorties and alert are possible given current aircrew resources, no analysis of overall wing capability is complete without considering the effects on maintenance of increasing output.

Major Results of Changes to Aircraft Scheduling

Rule and policy changes were made to the DOSS baseline model for aircraft scheduling that primarily increased aircraft availability to fly (such as increasing inspection intervals and turnaround sorties). Like the analysis of changes made to aircrew scheduling, the effects of changes to aircraft scheduling were analyzed at increased sortie and alert levels.

The major finding of the analysis was that aircraft availability is not a limiting factor in scheduling more sorties and alert. DOSS scheduled both increased sorties as well as increased aircraft on

alert, without the need of any policy changes to increase aircraft availability. These increases in output, however, necessitate scheduling more turnaround sorties, which allow a fixed number of aircraft to fly more sorties. Furthermore, increased numbers of turnaround sorties result in less aircraft hours spent in preflight and BPO inspections, which consequently reduces maintenance costs in terms of resources expended in these activities.

It should be stressed, however, that more turnarounds and fewer preflight and BPO inspections may have adverse effects on aircraft reliability over time. Other means of analysis besides DOSS would have to be employed to estimate the effects on aircraft reliability that may result from changes in aircraft scheduling.

Besides using DOSS to estimate the extent of sorties and alert that is possible given current resource levels, DOSS may be useful in the analysis of many other SAC resource allocation issues. In the following subsection I discuss the potential of DOSS to aid in the analysis of some of these other issues.

IMPLICATIONS OF DOSS FOR IMPROVING RESOURCE ALLOCATION WITHIN SAC

This study has demonstrated that DOSS provides a valid model of a typical wing's scheduling rules and policies. Furthermore, I have used DOSS as an analysis tool to examine particular resource allocation issues within SAC. My main purpose in this subsection is to discuss a broad range of other resource allocation issues within SAC where DOSS may be useful as an analysis tool. I also discuss some implications, not yet fully explored, for improving resource allocation within SAC by using DOSS in other ways, such as (a) a

training device for wing schedulers, and (b) as an aid in wing-level scheduling.

Other SAC Resource Allocation Issues

The analysis undertaken in this study focused on the issues of (a) flying more training sorties, and (b) increasing the alert force, while holding all resources constant. However, because DOSS incorporates resource constraints, Hq SAC policies, and wing-level decision rules, DOSS may have potential for examining a broad range of other SAC resource allocation issues as well.

For instance, without holding flying hours constant, we could use DOSS to examine the total number and mix of sorties that are possible at a typical SAC wing, given constraints only on aircrew and aircraft availability. Similarly, given rapidly rising fuel costs and continued budgetary cutbacks in SAC flying hours, we could use DOSS to examine alternative rules and policies under conditions of reduced flying hours.

The analysis undertaken in this study examined the extent to which the alert force could be increased, holding constant the number of aircrews and aircraft. However, DOSS has potential for examining a whole range of alert options at a variety of resource levels. For example, if the number of aircrews and aircraft are reduced a certain amount, what effect does this reduction have on alert capability given current scheduling rules and policies? To what extent can alternative rules and policies increase alert capability?

Furthermore, the rule-based methodology described in this study may have applicability in the analysis of different basing

alternatives. Clearly, basing issues involve many important concerns. The point here is that DOSS may have potential for shedding some light on parts of these very broad issues by incorporating the decision rules and policies for allocating resources. For example, one alternative to the current SAC basing structure is to consolidate B-52s and KC-135s at a few central locations, with satellite alert facilities dispersed throughout the CONUS (continental U.S.). We potentially could use DOSS to examine the effects of such a policy option on various measures of training, maintenance, and alert, that result from alternative rules and policies for resource allocation.

With the possible introduction of the air-launched cruise missile to the B-52 arsenal, and possibly subsequent changes in B-52 training and alert commitments, a host of issues arise involving tradeoffs associated with resource levels, sortie levels, maintenance, and alert. DOSS, as we have seen, is particularly well suited for examining these types of tradeoffs under alternative scheduling rules and policies.

The analysis undertaken in this study has concentrated solely on issues pertaining to the training, maintenance, and alert of B-52 aircrews and aircraft. Similar issues also exist pertaining to resource allocation of KC-135 aircrews and aircraft for which DOSS may be applicable. Moreover, resource allocation of KC-135 tankers is further complicated by the refueling responsibilities of these tankers, not only to SAC but to other Air Force Commands as well.

So far the kinds of issues which I have discussed, and the analysis which I have presented, involve the use of DOSS as an analysis tool at the SAC higher-headquarters level. I used DOSS to

analyze alternative rules and policies that affect resource allocation at SAC wings, in order to address policy issues of primary concern to higher-headquarters decisionmakers; namely, the extent of additional capability at SAC wings to fly more sorties, to increase alert, and to provide differential training. However, there are additional ways to use DOSS that may have potential implications for improving resource allocation within SAC.

Other Uses of DOSS for Improving Resource Allocation

Besides having potential as an analytic tool at the Hq SAC level, DOSS has implications not yet fully explored, as (a) a training device for wing schedulers, and (b) as an aid in wing-level scheduling.

DOSS as a Training Device. A major problem SAC wing schedulers face is the inability to see quickly what are the effects of a decision rule or particular set of decision rules on wing output. The relationships between decision rules and outcomes are simply too complex. This complexity compounds the problems of training new schedulers. Currently, there is no effective way to communicate the likely consequences of alternative decision rules on scheduling outcomes. New schedulers essentially learn by trial and error, through a process of on-the-job training.

DOSS, however, can incorporate a great deal of knowledge about scheduling in the form of decision rules employed by experienced schedulers. Furthermore, DOSS clarifies cause-and-effect relationships by allowing the user to vary decision rules and to see the consequences of such changes on schedules. If DOSS can effectively communicate this knowledge about decision rules and their

relationships to output, then DOSS has potential for training new or even experienced schedulers regarding the complexities of the scheduling process.

Although the potential of DOSS as a training device has yet to be systematically examined, a useful training aspect of DOSS has been observed in the course of this current research. The schedulers who participated in the development of the DOSS baseline models all indicated they gained valuable insights about the scheduling process from having to explicitly define their rules and rule interrelationships in helping to develop the models. It seems that by merely inputting to DOSS the necessary rules in order to adequately simulate scheduling, a great deal can be learned about the interrelationships of resource constraints, Hq SAC policies, and scheduling rules. An increased awareness of these interrelationships by schedulers could lead to better scheduling rules, and thereby better performance.

A means for more fully examining the potential of DOSS as a training device could be a useful by-product of the development of the system as an analytic tool at the Hq SAC level. Given an existing DOSS capability at the headquarters level, schedulers could be brought periodically to Hq SAC, as is currently done during quarterly scheduling conferences, to "experiment" with the system. That is, they could ask the system "what if" type questions varying scheduling rules that are within their discretion, and see the effects of alternatives on wing output. These exercises, besides providing training value, would also provide exposure of the kinds of wing-level scheduling problems DOSS can aid the scheduler in.

DOSS as a Wing-Level Scheduling Aid. Although this study has focused on the use of DOSS as an analysis tool at the Hq SAC level, wing schedulers could also use DOSS in a similar analytic fashion at the wing level. That is, by holding all Hq SAC policies and resource constraints constant, schedulers could use DOSS to examine the effects on scheduling outcomes of rules and parameters that are within their discretion. The use of DOSS in this fashion, essentially as a wing-level scheduling aid, has many implications, not yet fully explored, for improving resource allocation at the wing level.

For instance, DOSS facilitates the search for better schedules, and thereby has potential for increasing search and improving performance. As discussed earlier, because of the complexity of scheduling and cognitive limitations, schedulers follow the common decisionmaking process of "satisficing" as opposed to "optimizing." Because of the lack of clear cause-and-effect relationships, the enormous numbers of rules, and the vague nature of the goals, it is difficult enough for schedulers to develop even one "satisfactory" schedule, much less search for "better" ones. Because DOSS can rapidly generate different schedules, and provide operationally useful performance measures for making judgments about schedules, DOSS makes it considerably easier to search for better schedules.

Furthermore, DOSS can schedule many months of activity almost instantaneously, which could lead to the identification of potential problems sooner than would be expected if schedulers did not have such a scheduling aid. The rapid identification of a problem (such as not being able to meet an alert commitment three weeks from now) would then initiate a need for examining alternative solutions and, most important, allow enough time for a good solution to be found.

The use of DOSS as a wing-level scheduling aid may also lead to a reduction in slack resources paid to avoid uncertainty. As mentioned earlier, decisionmakers in organizations are quite reluctant to base activities on estimates of an uncertain future. They avoid uncertainty with payments of excess resources or slack. Slack resources serve as a hedge against uncertainty. For example, keeping an aircraft on the ground long enough to ensure (100 percent of the time) the completion of a desired task, is an expenditure of slack resources in terms of lowered aircraft availability to fly. Of course, the presence of slack is not necessarily inefficient. After all, some insurance may be worth the cost. However, without a clear understanding of the relationship of slack resources to output, it is difficult to strike the proper balance. In this regard, DOSS has potential to clarify the effects of slack on output. Schedulers can use DOSS to generate alternative schedules under different levels of slack, and can then examine the effects on a number of output measures. If overall performance could be shown improved with less slack, then this may lead to less attention being paid to uncertainty avoidance.

In addition, the use of DOSS at the wing level could improve coordination between the operations and maintenance suborganizations. As described earlier, operations and maintenance focus their attention on different goals, and consequently their preferred flying schedules rarely coincide. For example, takeoff times may be different, and the total sorties available according to maintenance may be less than originally requested by operations. Coordination eventually takes

place through an informal bargaining process, where both sides make adjustments from their initial positions. Because the effects of alternatives on overall wing performance are unclear, the bargaining outcome will not necessarily lead to the best alternative for improving overall wing performance. Other factors such as power struggles and gamesmanship largely determine bargaining outcomes (March and Simon, 1958:130), as opposed to some sort of cost and benefit analysis of alternatives.

DOSS has potential for improving the coordination between operations and maintenance because it can depict the effects of different flying schedules on output measures of concern to both operations and maintenance. For example, DOSS can depict the effects of alternative takeoff times on training output, as well as on the workload patterns of maintenance crews. Therefore, a scheduler, whether he is in maintenance or operations, could use DOSS to argue his case by showing how his suggestions compare to others in terms of overall wing outcomes. Potentially, using DOSS in this fashion could lead to a coordination process that relies more on the analysis of alternatives in terms of overall wing performance, as opposed to being dominated by bargaining strategies. Conceivably, DOSS could even make a form of centralized wing scheduling possible, where operations and maintenance goals are considered simultaneously. For example, with decision support from DOSS, aircrews and aircraft could be assigned simultaneously to training sorties. Any particular assignment could be made by weighing the benefits in terms of training gained by the aircrew against the cost to maintenance in supplying the aircraft.

Future Work

Given the potential implications of DOSS for improving resource allocation within SAC, many issues remain requiring additional investigation as to how best to use such a technology. This study has attempted to demonstrate the potential of DOSS as an analytic tool to address various types of resource allocation policy issues of particular concern to Hq SAC decisionmakers. Much more work remains in terms of further defining the demand for such a methodology at the Hq SAC level--that is, to determine the extent of issues facing SAC now and in the future where DOSS would provide a useful methodology for analysis. Furthermore, there need to be concurrent investigations regarding the feasibility and costs of adapting the current DOSS prototype technology to alternative computer facilities available to the Hq SAC staff. Currently, SAC is considering these issues of demand and cost of the DOSS prototype system for use as an analysis tool at the headquarters level.

If on the basis of future work DOSS was developed for analysis purposes at the headquarters level, such a development would provide a means for evaluating other potential uses of DOSS. That is, a DOSS capability at the Hq SAC level would provide a means for evaluating the potential of DOSS as a training device for wing schedulers. This in turn would provide continued exposure of DOSS to wing schedulers, which would be useful for evaluating potential ways DOSS could aid the scheduler at the wing level.

IMPLICATIONS OF RULE-BASED MODELING FOR IMPROVING ORGANIZATIONAL
DECISIONMAKING

The study has focused on the potential of a rule-based model (DOSS) for improving resource allocation within a large public-sector organization (SAC). My intent in this subsection is to discuss some implications, stemming from this research, regarding the potential of rule-based modeling for improving aspects of organizational decisionmaking in general. Therefore, the discussion that follows is no longer bound to the issues surrounding the development and use of DOSS for SAC, or even for other Air Force Commands. However, it should be stressed that only through further research, in a variety of contexts, can we thoroughly validate any generalizations from the findings of this study.

In general, rule-based models have potential implications for improving organizational decisionmaking through their use as (1) descriptive models that simulate decisionmaking within organizations, and (2) as analytic tools for examining the effects of alternatives.

Implications as Descriptive Models

The fundamental idea underlying rule-based modeling is that in order to improve decisionmaking within organizations you first have to understand it. In this regard, the theory of decisionmaking that served as the basis for the design and use of DOSS is relevant for the design and use of rule-based models in general. This body of theory evolved largely from theoretical studies of organizational decisionmaking done at the Carnegie Institute of Technology during the late 1950s and early 1960s. The works of Simon (1957a), March and

Simon (1958), and Cyert and March (1963), have been instrumental in developing a general theory of decisionmaking within organizations based on the notion of "cognitive limits on rationality." Simply put, in order to deal with complexity, decisionmakers rely on many heuristics or rules of thumb for making decisions. Furthermore, they search in a limited fashion for solutions that are "good enough."

This school of thought, which Steinbruner (1974) refers to as the cybernetic paradigm, was largely formulated and tested in the context of private-sector firms. However, because the paradigm relies on basic cognitive principles, its concepts are not intrinsically unique to the business firm. Such concepts as factorization, standard operating procedures, uncertainty avoidance, and satisficing apply equally as well to decisionmaking in public-sector organizations. For example, Wildavsky's (1964) theory of the governmental budgetary process fits quite well within the cybernetic paradigm. In addition, Crecine (1969) has observed aspects of the paradigm with regards to municipal resource allocation. Allison (1971) describes how many aspects of the paradigm can be applied to explain the actions of various governmental organizations involved in the Cuban Missile Crisis. In fact, one purpose of this study has been to provide further evidence in support of the cybernetic paradigm with respect to internal resource allocation in large military organizations such as SAC.

Because of the rule-based nature of cybernetic decision processes, and their prevalence in many organizational settings, rule-based modeling holds promise as a tool for building and testing the validity of descriptive models of decisionmaking in a variety of

organizations. One of the earliest rule-based models, developed by Clarkson (1962), simulated the rules for portfolio selection by trust officers of a large bank. Cyert and March (1963) programmed into a computer the decision rules which they theorized reflected the price and output decisionmaking behavior of a private-sector business firm. They then tested their theories by comparing the outputs of their model with actual decisions made by the firm. Crecine (1969) used a similar rule-based computer approach to test his theories of the decisionmaking process prevalent in municipal resource allocation.

Besides the advantage of theory formalization and testing cited above, a descriptive rule-based model of the decisionmaking processes of an organization also has certain normative implications. For example, Cyert and March (1963:290) have suggested that the description of an organization in terms of its decisionmaking processes, as opposed to more standard descriptions such as organizational charts, may stimulate consideration of a different set of alternative changes in the system. Furthermore, as was observed in the development of DOSS, the process of developing a descriptive rule-based model of an organization may lead to valuable insights regarding the complexities of the decisionmaking process itself. However, although a descriptive rule-based model abstracts important characteristics of organizational decisionmaking, it alone does not directly point to better solutions. In order to more directly improve decisionmaking, rule-based models must be used in an analytic fashion to explore the consequences of alternatives.

Implications as Analysis Tools

There are two fundamental ways in which rule-based models have potential for analysis, both of which have different implications for improving organizational decisionmaking. Rule-based models can be used for (1) analyzing decisionmaking within an organization in order to seek improvements, and (2) analyzing decisionmaking of outside organizations for the purposes of prediction and control.

Analyzing Decisionmaking within an Organization. Rule-based models that accurately simulate decisionmaking within an organization not only attempt to reproduce final output, but also the decision rules, corporate-level policies, and external constraints that influence final output. Therefore, rule-based models can be used to examine the effects of alternative decision rules, policies, and external constraints on organizational outcomes. DOSS was used in this fashion for analyzing resource allocation issues within SAC, and therefore provides some evidence in support of the generalizations that are made.

Because of the complexities of problems that confront large organizations, organizations typically factor their decision problems into subproblems and assign the subproblems to subunits within the organization. This factorization of decision problems generally leads to "local rationality," where subunits deal only with a limited set of problems and consequently a limited set of subgoals (Cyert and March, 1963:117). With regard to SAC resource allocation, for example, we have seen how the problems and goals of operations are quite different from those of maintenance.

Because of factorization and local rationality, rule-based models of decisionmaking within subunits of an organization have different implications depending on how they are used. For instance, we saw and discussed how DOSS could be used not only as an aid to wing-level scheduling, but also as a tool for analyzing alternative resource allocation policy issues of concern to Hq SAC decisionmakers. In general, rule-based models of organizational subunits have implications not only for improving operational decisionmaking within the subunits, but also for improving corporate-level decisionmaking (i.e., decisionmaking at the central or headquarters level of an organization).

At the subunit level, a rule-based model could help analyze the effects of alternative decision rules on the performance of the subunit, holding constant corporate-level policies and external constraints. Such an analysis tool has several implications for improving decisionmaking with respect to the subunit's goals, such as: (1) increased search for better alternatives, (2) more rapid recognition of potential problems with current decision rules, and (3) possible reductions in excess resources (slack) paid to avoid uncertainty.

Rule-based modeling may increase the search for better alternatives. As we have seen with regard to SAC resource allocation, and as others have observed in different organizational contexts, the search for alternative solutions is greatly limited (Simon, 1957a; March and Simon, 1958; and Cyert and March, 1963). Due to cognitive limits and the lack of clear cause-and-effect relationships, search does not cover all possible alternatives in the hope of finding the

optimal solution. Instead, search is based on a limited set of rules, and usually stops with the discovery of a feasible or satisfactory solution. Rule-based models, however, facilitate the search of alternative decision rules for solving problems, and thereby may lead to increased search. Rule-based models attempt to expand cognitive limits and help clarify cause-and-effect relationships. They do not seek the elusive optimal solution, but rather aim only to depict the likely consequences of alternatives.

Rule-based modeling may also lead to more rapid recognition of potential problems. In most organizations search is not only limited, but it is also problem-directed (Cyert and March 1963:121). That is, search is initiated only when current decision rules no longer lead to satisfactory outcomes. Furthermore, a prime instigator of problems for an organization is a changing environment, caused largely by changes in other organizations such as competitors or governments. Moreover, problems with current decision rules are often not recognized in time to find good solutions. Rule-based modeling, however, could lead to more rapid recognition of potential problems with current decision rules. For instance, rule-based models could be used to examine the effects of current decision rules given changes in external constraints. Recognition of potential problems could lead to search for solutions, well in advance of the time the problems with the current rules would normally become evident or extremely costly to overcome.

In addition, rule-based modeling may lead to a reduction of slack resources paid to avoid uncertainty. A widespread feature of organizational decisionmaking, and one we have observed within the SAC

wing environment, is that decisionmakers tend to avoid uncertainty (Cyert and March, 1963:119). Thompson (1967:159) goes so far as to say, "Uncertainty appears as the fundamental problem for complex organizations, and coping with uncertainty, as the essence of the administrative process." The commitment to attaining certainty leads, among other things, to the payment of excess resources or slack to act essentially as insurance against uncertain outcomes (Cyert and March, 1963:118-120; Thompson, 1967:150; Downs, 1968:138). The existence of some slack, however, is not necessarily inefficient. As mentioned earlier, some insurance may be worth the cost. However, without a clear understanding of the relationship of slack to output it is difficult to strike the proper balance. Rule-based models, however, have the potential for clarifying the effect of alternative levels of slack. Rule-based models allow decisionmakers to examine the effects of different decision rules that reflect different levels of slack resources. Given that increased performance can be shown possible as a result of less slack, decisionmakers are then made aware of the opportunity costs of slack, and are thus in a better position to make more informed choices.

Although rule-based modeling has several implications for improving decisionmaking at the subunit level, decisionmaking will still be driven by subunit goals, and not necessarily by overall organizational goals. It is at the corporate level of decisionmaking (i.e., the central or headquarters level) where the main responsibility lies for coordinating decisionmaking across subunits for the purposes of achieving overall organizational goals. It is at the corporate level where tradeoffs between conflicting subunit goals

need to be made. Because of the extreme complexity of finding, or even defining, the "optimal" solution with respect to overall organizational goals (which must take into account all subunit subgoals as well as the development of the "optimal" level of organizational slack), rule-based modeling may have its most significant potential with regard to aiding corporate-level decisionmaking.

The use of rule-based models at the corporate level has implications for improving overall organizational performance by (1) allowing increased coordination of decisionmaking across subunits, and by (2) aiding in the analysis of corporate policy alternatives.

Rule-based modeling may allow increased coordination across subunits. As mentioned earlier, complex problems are factored out to subunits who deal only with a limited set of subgoals. Depending upon the degree of interdependence among subunits, different subgoals often lead to intergroup conflicts within the organization (March and Simon, 1958:121-129). March and Simon discuss two basic processes by which intergroup conflicts can be resolved: the analytic and bargaining processes (1958:130). The analytic process basically involves considering the different actions of subunits in terms of the effects on overall organizational outcomes. The bargaining process is characterized by little consideration of overall organizational outcomes, but rather involves each subunit trying to better their own position in terms of their own subgoals. Because of cognitive limitations, and the lack of clear cause-and-effect relationships in complex organizational systems, resolving conflict through analysis is extremely burdensome. Intergroup conflict within organizations is

usually resolved through bargaining, where slack resources provide a means to absorb potential conflicts (March and Simon, 1958:126; Cyert and March, 1963:118), and where organizational outcomes are largely determined by power struggles and gamesmanship (March and Simon, 1958:130).

Rule-based models help clarify cause-and-effect relationships and help expand cognitive limitations. Therefore, rule-based models may aid in the analysis of alternative actions of subunits in terms of their effects on the organization as a whole. The use of such an analysis tool at some central level of decisionmaking, where the responsibility rests for making tradeoffs among subgoals, could facilitate resolving conflicts through a form of analysis, and thereby reduce reliance on bargaining processes. For instance, I have discussed the possible use of DOSS in this regard for resolving conflicts between operations and maintenance at SAC bomb wings.

Rule-based models may also aid in the analysis of corporate-level policy alternatives. Rule-based models of decisionmaking within subunits of an organization attempt to incorporate not only operational decision rules, but also corporate-level policies and external constraints that influence final output. Therefore, corporate-level analysts or planners could use rule-based models to examine the effects of alternative corporate-level policies, holding all other factors constant. In other words, rule-based models could act as simulation models to predict the likely effects on subunit performance that would result from different corporate policies. In addition, corporate-level analysts or planners could use rule-based models to examine the changes to decision rules that would be

necessary in order to implement different corporate policies. For instance, I used DOSS in this regard for estimating the extent of additional capability at SAC bomb wings to fly more sorties and increase the alert force.

In addition to implications for improving decisionmaking at different levels within an organization, rule-based modeling also has implications for analyzing decisionmaking of outside organizations.

Predicting and Controlling the Behavior of Other Organizations.

Decisionmaking within an organization is affected directly and indirectly by a variety of other complex organizations in the external environment. Obviously, decisionmaking within most organizations could be improved by better knowledge of the behavior of competitors, suppliers, governmental organizations, labor unions, and many other organizations. Rule-based models that can predict the decisionmaking behavior of these other organizations clearly have implications for improving decisionmaking. As Cyert and March point out with regard to their model of decisionmaking of a retail department store, ". . . if we can predict the ordering behavior of a retail department store, we should be able to provide useful information to a manufacturer selling to such a store (1963:291)." Basically, rule-based models of other organizations could be used to help analyze decision alternatives in terms of their effects on other organizations.

Rule-based models of the internal decision processes of other organizations could also be useful for analyzing external policies that are designed to control aspects of organizational behavior. This use of rule-based modeling is particularly relevant for most governmental regulatory agencies, who seek to influence organizational

decisionmaking for some societal purpose. Whether public policies, such as economic or environmental regulations, will lead to their desired ends depends largely upon how they enter the internal decision processes of the organizations they are directed toward. Rule-based models, to the extent they can help clarify how an organization will deal internally with policy variables being modified externally, have implications for analyzing the likely effects of alternative policies.

To be useful in predicting the behavior of other organizations, however, rule-based models have a serious limitation: they require a great deal of detailed knowledge about the decision processes and internal operations of those organizations, which may be difficult to obtain. Few organizations will knowingly provide detailed information to their competitors, or to governmental organizations interested in regulating them. Consequently, much more research is needed to either (1) develop more general theories of decisionmaking that limit the number of context-dependent characteristics that have to be incorporated, or (2) develop a great many rule-based models of decisionmaking in a variety of organizational contexts.

In conclusion, I have drawn on experience with DOSS to suggest ways in which rule-based modeling might improve organizational decisionmaking. Much more research is needed, however, to validate some of the generalizations that were made. To my knowledge, no direct evidence exists regarding the potential of rule-based modeling for analyzing the effects of external public policies on organizational decisionmaking. Nevertheless, a better understanding of how public policies influence decisionmaking within organizations is of paramount importance for developing effective public policy.

This is further reason why rule-based modeling as an analysis tool of organizational decisionmaking seems worthy of a great deal more research in a variety of organizational contexts.

Appendix A

TRAINING ITEMS

This appendix describes the method used to estimate total training items, as well as their distribution, that result from the alternative schedules examined in this study.

In order to estimate the total numbers of training items typically scheduled for different types of sorties, data were collected from weekly schedules of Wing A for the period September through November 1978. This data source records in detail what training items are scheduled for each sortie. The total number of six important training items that are typically scheduled on each type of sortie was computed. These six items, which reflect the major training component of SAC training sorties, are the following:

- o B01s -- Low Altitude Bomb Runs (bombing exercises)
- o N09s -- TA/EVS Navigation Legs (terrain avoidance)
- o N15s -- Low Altitude Navigation Legs (navigation)
- o R01s -- Air Refueling
- o E01s -- RBS ECM Runs (electronic countermeasures)
- o F01s -- Fighter Intercepts (exercises with fighter aircraft).

Sorties were classified as either non-CSS event sorties, CSS sorties, day profile sorties, or night profile sorties. Pilot pro sorties were not considered since these sorties are primarily for training pilots in instrument approaches and emergency procedures.

The average number of training items scheduled for a particular type of sortie is reported below in Table A1. For example, at Wing A 136 Low Altitude Bomb Runs (B01s) were scheduled on 76 non-CSS event training sorties, which resulted in an average of 1.79 B01s for each scheduled non-CSS event sortie.

Table A1
TRAINING ITEMS SCHEDULED PER SORTIE

Sortie Types Scheduled	B01s	N09s	N15s	R01s	E01s	F01s
Non-CSS Event Sorties N=76	1.79	.89	.89	.48	1.79	.20
CSS Sorties N=46	3.20	1.0	1.0	1.0	2.20	.00
Night Profile Sorties N=57	2.95	1.0	1.0	1.0	2.95	.04
Day Profile Sorties N=47	1.98	1.0	1.0	1.0	1.98	.19

SOURCE: Weekly schedules of Wing A.

Different mixes of scheduled sortie types naturally result in different numbers of training items scheduled, as shown in Table 13. In order to estimate the number of training items scheduled, given a particular sortie mix (as was done for Table 13), I simply multiplied the number of scheduled sorties of different types by their average training content per scheduled sortie. For example, if 100 non-CSS event sorties were scheduled, I would estimate that approximately 179 B01s would be scheduled (100 sorties x 1.79 B01s per sortie).

In addition, different sortie distributions allocated to MR crews naturally lead to different training item distributions. As is observed in Section V, widely varied distributions of sorties lead to widely varied training item distributions. Furthermore, differential distributions of sorties based on experience levels lead to differential distributions of training items.

To estimate the number of training items scheduled to each crew, in order to compute expected training item distributions, I simply multiplied the number of sorties of different types scheduled to a particular crew by the average training content per scheduled sortie. For example, if a crew was scheduled five non-CSS event sorties I would estimate that crew would be scheduled approximately 9 B01s ($5 \text{ sorties} \times 1.79 \text{ B01s per sortie}$) on those five sorties.

Appendix B

DISTRIBUTION OF SORTIES

This study involved two goals for distributing sorties to mission-ready (MR) crews. The first goal was to give equal numbers of sorties to all MR crews. To examine the extent this was achieved under alternative conditions, the ranges and standard deviations of sortie distributions were compared. The second goal for distributing sorties was to provide differential training based on relative experience levels. Thus, "S" crews, "E" crews, and "R" crews were considered to have different total sortie requirements, as well as different requirements for sorties of different training content (see Table 5).

To examine the effect of differential sortie requirements on sorties scheduled, Ordinary Least Squares (OLS) regressions were calculated where the independent variable was "total sortie requirements" and the dependent variable was "total sorties scheduled." A desirable allocation, in terms of differential training, would have sorties scheduled equal to sorties required, with any excess sorties above minimum requirements distributed equally. Such an allocation would be represented by the linear model depicted below:

$$y = b_0 + b_1 x$$

where

- y = total sorties scheduled
- x = total sortie requirements
- b₀ = a constant term reflecting the average number of excess sorties above minimum requirements
- b₁ = 1.0 (i.e., the slope of the OLS line)

The above relationship is the objective that DOSS sought in attempting to provide differential training. Naturally, many other variables such as aircrew availability affect how sorties are scheduled. Thus, we expect to see some variance in sorties scheduled not explained by sortie requirements.

The regression coefficients, R-squared statistics, and t-statistics for the regressions involving total sorties are reported below in Table B1. One separate regression is reported for each of the alternative schedules that were examined in this study. Given that each regression involves 15 degrees of freedom, a t-statistic value greater than 2.13 is statistically significant at the .05 significance level.

Table B1

RESULTS OF REGRESSIONS ON TOTAL SORTIES SCHEDULED

Alternative Schedules	Constant Term	Slope	R-Squared	t-Statistic
Baseline Sortie & Alert Level: -----				
Wing (a)	12.42	.09	.02	.50
DOSS-Differen- tial Training	7.59	.48	.44	3.41
DOSS-Differen- tial Training & Increased Aircraft Availability	3.63	.79	.77	6.70
Baseline Sortie Level with Increased Alert: -----				
DOSS-Differen- tial Training	10.27	.20	.08	1.15
DOSS-Differen- tial Training & Increased Aircraft Availability	2.80	.86	.86	9.60
Increased Sorties and Alert: -----				
DOSS-Differen- tial Training & Increased Aircraft Availability	1.40	.89	.82	8.30

^a
Wing schedules of Wing A.

To examine the effect of differential requirements for sorties of different training content, I analyzed the distributions of non-CSS event sorties in detail. OLS regressions were calculated where the independent variable was "non-CSS event sortie requirements" and the dependent variable was "non-CSS event sorties scheduled." As in the case of total sortie distributions, a desirable allocation in terms of differential training is reflected by the slope of the estimated OLS line being nearly equal to 1.0. The regression coefficients, R-squared statistics, and t-statistics for these regressions are reported below in Table B2. Again, one separate regression is reported for each of the alternative schedules that were examined in this study. Also, given that each regression involves 15 degrees of freedom, a t-statistic value greater than 2.13 is significant at the .05 level.

Table B2

RESULTS OF REGRESSIONS ON NON-CSS EVENT SORTIES SCHEDULED

Alternative Schedules	Constant Term	Slope	R-Squared	t-Statistic
Baseline Sortie & Alert Level: -----				
Wing (a)	4.2	.02	.0004	.03
DOSS-Differential Training	2.3	.52	.34	2.8
DOSS-Differential Training & Increased Aircrew Availability	1.1	.79	.61	4.8
Baseline Sortie Level with Increased Alert: -----				
DOSS-Differential Training	2.6	.42	.17	1.6
DOSS-Differential Training & Increased Aircrew Availability	1.8	.65	.32	2.7
Increased Sorties and Alert: -----				
DOSS-Differential Training & Increased Aircrew Availability	.6	.99	.64	5.2

^a
Wing schedules of Wing A.

GLOSSARY

BO1	Low Altitude Multiple Release Bomb Run--a bombing exercise common in most SAC training sorties.
BPO	Basic Postflight inspection following most sorties.
CCRR	Combat Crew Rest and Recuperation--a specified period of crew rest following an alert duty.
CSS	Cold Seat Swap sortie--a sortie flown by an aircraft which has just landed from a previous sortie, where virtually no maintenance is performed in between.
DOSS	Decision Oriented Scheduling System
DNIF	Duty Not to Include Flying--caused usually by illness or injury.
"E" Crew	Senior Crew--middle level of aircrew proficiency
ECM	Electronic Countermeasures
EPTS	EWO Profile Training Sortie--an integral crew sortie that simulates the EWO mission.
ETS	Event Training Sortie--an integral crew sortie designed to improve crew proficiency in low-level penetration and bombing.
EWO	Emergency War Orders--the type of mission toward which SAC training is directed.
4X	A sortie where less than a specified number of crew members are from the same crew.
HHDM	Higher Headquarters Directed Mission--a requirement for the generation of a sortie, levied by an authority above wing level.
Hq SAC	Headquarters, Strategic Air Command

MR	Mission-Ready--a crew or crew member certified as ready for combat.
OLS	Ordinary Least Squares regression--a statistical technique.
PDM	Programmed Depot Maintenance--a requirement that each aircraft in the fleet periodically go to an Air Force depot for depot-level maintenance.
Pilot Pro	Pilot Proficiency Sortie--a sortie designed to concentrate on maintaining pilot skills only.
Pre-flight	Basic preflight inspection required before most sorties.
PTS	Profile Training Sortie--an integral crew sortie requiring low-level navigation, low-level bombing, and air refueling exercises.
"R" crew	Ready Crew--lowest level of aircrew proficiency.
RBS site	Radar Bomb Scoring Site--a facility capable of scoring simulated bombing and producing various electronic jamming signals.
"S" crew	Selected Crew--highest level of aircrew proficiency.
SAC	Strategic Air Command
SACM	SAC Manual--a Hq SAC manual that is directive in nature to subordinate units.
SACR	SAC Regulation--a manual stipulating Hq SAC regulations.
Stand/Eval	At each wing three of the most qualified mission-ready crews are designated to periodically evaluate other crews in flight.
TA	Terrain Avoidance--procedures for flying at low altitudes.

TCTO	Time Compliance Technical Order--a time-urgent modification to be performed on an aircraft.
TDY	Temporary duty of short duration, away from the home station.
Thru-flight	An inspection usually made to an aircraft prior to flying a turnaround sortie.
Turnaround	A sortie where the aircraft flown does not require a preflight inspection, because its last preflight inspection has been within 72 hours.

BIBLIOGRAPHY

- Allison, Graham (1971): Essence of Decision, Little, Brown and Company, Inc., Boston.
- Becker, G. B., and C. G. McClintock (1967): "Value: Behavioral Decision Theory," Annual Review of Psychology, Vol. 18, pp. 239-286.
- Berman, Morton B. (1974): Improving SAC Aircrew and Aircraft Scheduling to Increase Resource Effectiveness, The Rand Corporation, R-1435, Santa Monica, California.
- _____, (1975): Scheduling Aircrews and Aircraft: Problems of Resource Allocations in the Strategic Air Command, The Rand Corporation, R-1610, Santa Monica, California.
- Buchanan, J. R., and R. D. Fennel (1977): "An Intelligent Information System for Criminal Case Management in the Federal Courts," Proc. 5th Int. Joint Conf. Artificial Intelligence, Cambridge, Massachusetts, pp. 901-902.
- Clarkson, G.P.E. (1962): Portfolio Selection: A Simulation of Trust Investment, Prentice-Hall, Inc., Englewood Cliffs, New Jersey.
- Cohen, I. K. (1972): Aircraft Planned Inspection Policies: A Briefing, The Rand Corporation, R-1025-PR, Santa Monica, California.
- Crecine, John P. (1969): Governmental Problem-Solving, Rand McNally, Chicago.
- Cyert, Richard M., and James G. March (1963): A Behavioral Theory of the Firm, Prentice-Hall, Inc., Englewood Cliffs, New Jersey.
- Downs, Anthony (1967): Inside Bureaucracy, Little, Brown and Company, Inc., Boston.
- Edwards, Ward (1961): "Behavioral Decision Theory," Annual Review of Psychology, Vol. 12, pp. 473-498.
- Elwell, Ralph, and Chris Roach (1976): Scheduled Maintenance Policies for the F-4 Aircraft: Results of the Maintenance Posture Improvement Program, The Rand Corporation, R-1942-PR, Santa Monica, California.
- Fallon, Richard (1979): DOSS as an Analysis Tool: Applications to the B-52 Aircrew Continuation Training Test Program, The Rand Corporation, N-1161-1-AF, Santa Monica, California.

- Feigenbaum, E. A. (1977): "The Art of Artificial Intelligence: Themes and Case Studies of Knowledge Engineering," Proc. 5th Int. Joint Conf. Artificial Intelligence, Cambridge, Massachusetts, pp. 1014-1029.
- Feigenbaum, E. A., B. G. Buchanan, and J. Lederberg (1971): "On Generality and Problem Solving: A Case Study Using the DENDRAL Program," in B. Meltzer and D. Michie (eds.), Machine Intelligence, Vol. 6, American Elsevier, New York, pp. 165-190.
- Johnson, Erik (1968): "Studies in Multiobjective Decision Models," The Economic Research Center in Lund, Lund, Sweden.
- Keen, P.G.W., and M. S. Scott Morton (1978): Decision Support Systems: An Organizational Perspective, Addison-Wesley Publishing Company, Inc., Reading, Mass.
- Leger, Carl R., and W. L. Peacock (1975): Logistics Scheduling Rules Study Data Collection Report, compiled for Air Force Systems Design Center by Potomac Research, Inc.
- Luce, R. D., and H. E. Raiffa (1957): Games and Decisions, John Wiley and Sons, New York.
- Luce, R. D., and P. Suppes (1965): "Preference, Utility, and Subjective Probability," in R. P. Luce, R. R. Bush, and E. Galanter (eds.), Handbook of Mathematical Psychology, John Wiley and Sons, New York.
- March, James G., and H. A. Simon (1958): Organizations, John Wiley and Sons, New York.
- Newell, Allan, and H. A. Simon (1958): "Heuristic Problem Solving: The Next Advance in Operations Research," Operations Research, Vol. 6, pp. 1-10.
- _____, (1972): Human Problem Solving, Prentice-Hall, Englewood Cliffs, New Jersey.
- Reitman, W. (1965): Cognition and Thought, John Wiley and Sons, New York.
- Shortliffe, E. H. (1976): Computer-Based Medical Consultations: MYCIN, American Elsevier Publishing Company, Inc., New York.
- Simon, Herbert A. (1957a): Administrative Behavior, 2d ed., Macmillan, New York.
- _____, (1957b): Models of Man: Social and Rational, John Wiley and Sons, New York.

- _____, (1959): "Theories of Decision Making in Economics," Am. Econ. Rev., Vol. 49, pp. 253-283.
- _____, (1960): The New Science of Management Decision, Harper and Row, Inc., New York.
- _____, (1966): "The Logic of Heuristic Decision Making," in N. Rescher (ed.), The Logic of Decision and Action, University of Pittsburgh Press, Pittsburgh, Pennsylvania.
- _____, (1968): The Sciences of the Artificial, MIT Press, Cambridge, Massachusetts.
- Steinbruner, John D. (1974): The Cybernetic Theory of Decision: New Dimensions of Political Analysis, Princeton University Press, Princeton, New Jersey.
- Strauch, Ralph E. (1974): A Critical Assessment of Quantitative Methodology as a Policy Analysis Tool, The Rand Corporation, P-5282, Santa Monica, California.
- Sylvan, D. A., E. Thorson, and S. J. Thorson (1977): Processing Styles and Theories of Foreign Policy Decision Making, International Studies Association, St. Louis.
- Thompson, James D. (1967): Organizations in Action, McGraw-Hill, New York.
- U.S. Congress, (1979a): Department of Defense Appropriation Bill, 1979, U.S. House of Representatives, 95th Cong., 2d sess., No. 95-1398.
- _____, (1979b): Department of Defense Appropriations for 1980, U.S. House of Representatives, 96th Cong., 1st sess., p. 397.
- Von Neumann, J., and O. Morgenstern (1954): Theory of Games and Economic Behavior, John Wiley and Sons, New York.
- Waterman, D. A., and B. Jenkins (1976): Heuristic Modeling Using Rule-Based Computer Systems, The Rand Corporation, P-5811, Santa Monica, California.
- Wesson R. B. (1977): "Planning in the World of the Air Traffic Controller," Proc. 5th Int. Joint Conf. Artificial Intelligence, Cambridge, Mass., pp. 473-479.
- Wildavsky, Aaron (1964): The Politics of the Budgetary Process, Little, Brown and Company, Inc., Boston.

LMED
-8